Research Article

Rational arbitration between statistics and rules in human sequence learning

Maxime Maheu[1,2,a]*, Florent Meyniel[1,b†] and Stanislas Dehaene[1,3,c†]

Abstract

Detecting and learning temporal regularities is essential to accurately predict the future. Past research indicates that humans are sensitive to two types of sequential regularities: deterministic rules, which afford sure predictions, and statistical biases, which govern the probabilities of individual items and their transitions. How does the human brain arbitrate between those two types? We used finger tracking to continuously monitor the online build-up of evidence, confidence, false alarms and changes-of-mind during sequence learning. All these aspects of behaviour conformed tightly to a hierarchical Bayesian inference model with distinct hypothesis spaces for statistics and rules, yet linked by a single probabilistic currency. Alternative models based either on a single statistical mechanism or on two non-commensurable systems were rejected. Our results indicate that a hierarchical Bayesian inference mechanism, capable of operating over several distinct hypothesis spaces, underlies the human capability to learn both statistics and rules.

[1]Cognitive Neuroimaging Unit, CEA DRF/JOLIOT, INSERM, Université Paris-Sud, Université Paris-Saclay, NeuroSpin centre, Gif-sur-Yvette, France. [2]Université de Paris, Paris, France. [3]Collège de France, Paris, France. [a]https://orcid.org/0000-0002-6851-4927, [b]http://orcid.org/0000-0002-6992-678X, [c]https://orcid.org/0000-0002-7418-8275
†These authors share senior authorship. *e-mail: maheu.mp@gmail.com

From weather to traffic lights, many real-life processes unfold across time, forming sequences of events that are characterised by some inner structure. The ability to learn such sequential regularities is essential to agents navigating real-life environments because it enables them to make predictions about the future [1–3]. Past research indicates that the brain constantly entertains predictions about the future [4–6] and leverages those predictions to promote a more efficient processing of incoming information as well as improve decision-making and behavioural control [7–9].

In order to make accurate predictions, human observers must solve two difficult problems. First, they should *detect when a regularity appears* [10–14]. Second, because various types of regularities exist, observers must also *identify the kind of process generating the regularity* in order to ensure accurate predictions. In theory, the number of hidden generative processes is infinite. Here, we study a simple distinction between statistical biases and deterministic rules, which relates to the strength of predictions they afford (*Fig. 1a*). Statistical biases allow uncertain, yet better-than-chance, predictions. For instance, looming dark clouds are generally followed by rain, but sun may also appear in rarer occasions. Deterministic rules, on the other hand, allow sure predictions. For instance, at a French traffic light, only the red light (stop signal) can follow the amber one (warning signal).

In our taxonomy, deterministic rules are an extreme case of statistical biases, in which the probability of the next event is 0 or 1 (*Fig. 1b*). We hypothesize that humans treat statistics and rules as fundamentally different, corresponding to two distinct hypothesis spaces. Those two types of regularity allow predictions of a different nature (uncertain vs. certain), they enable computations of different kinds (which we detail later), and in practice, they have been explored in mostly distinct lines of research. On the one hand, detecting, estimating or leveraging statistical biases is at the heart of studies on probability learning [14–21] and reinforcement learning [22–28]. On the other hand, deterministic rules were embedded in sequence learning and artificial grammar paradigms, in which repeated patterns have to be detected or reproduced [10,13,29–37]. Studies that combine or compare both kinds of regularities appear as exceptions [38–43].

In this article, we propose a new paradigm in order to jointly study the problems of detecting when a regularity emerges, and of identifying its type. This approach offers opportunities to bridge gaps between previous studies which investigated those aspects one at a time. Notably, this approach enables to test a new normative theory of rational inference, which assumes that the detection and identification of regularities should rely on a hierarchical Bayesian inference over a minimum of three distinct hypothesis spaces: fully random, statistical bias, and deterministic rule. This new proposal makes several predictions, in terms of subjects' performance in detecting regularities and discriminating between types of regularities, the strength of their beliefs, the specific dynamics associated with the detection, and even the type of errors people make, which we test in our task. Some of those predictions are unique to this proposal and support the idea that statistics and rules pertain to different hypothesis spaces, rather than a gradation along a continuum.
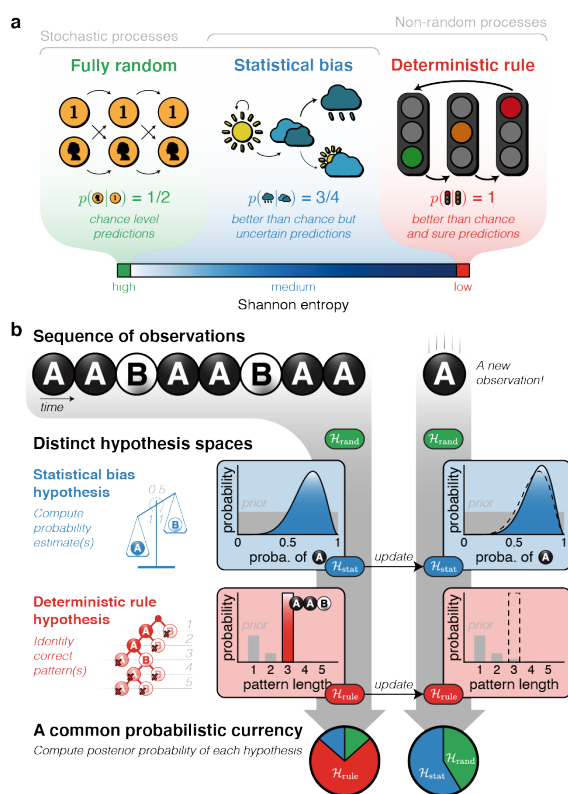

Fig. 1 | A taxonomy of regularities and a normative model of how those regularities can be inferred from a sequence of observations. **a**, *A taxonomy of regularities*. A first distinction separates random (i.e. unpredictable) from non-random (i.e. predictable) processes. A second distinction, with non-random processes, separates statistical biases (i.e. affording only uncertain predictions) from deterministic rules (i.e. affording certain predictions). Shannon entropy quantifies the unpredictability of a process and therefore differs between those different types of processes. **b**, *Rational inference from an example sequence*. An example sequence of binary observations is presented to the model which arbitrates between different hypothesis by means of probabilistic inference. For simplicity, we supposed here that there is no volatility: the full binary sequence corresponds to a single generative process, which can be either random, or exhibit a statistical bias (here, in the frequency of items) or follow a deterministic rule (here, the repetition of a fixed pattern of at most five items). Following the first 8 observations, the statistical bias hypothesis estimates than As are overall more likely than Bs. By contrast, the deterministic rule hypothesis estimates that the sequence can be described as the repetition of the AAB pattern. Computing the posterior probability of each hypothesis reveals that, at this point, the deterministic rule hypothesis is the most likely hypothesis. Note, however, that a single new observation (here, the final A) suffices to discard the deterministic rule hypothesis.

Results

Experimental design. To test our theoretical proposal, we designed a novel behavioural experiment in which we presented human subjects with binary auditory sequences (200 observations, for 1 minute each; *Fig. 2a*). All sequences started with a random process (i.e. like tossing a fair coin). In about two thirds of the cases, after a delay of variable length, we introduced either a statistical bias (in the first-order transition probabilities) or a deterministic rule (the repetition of a particular pattern of length-4 to 10). In the remaining cases, the sequences remained fully random until the end. Subjects were fully instructed about this task structure and were asked to slide their finger towards one of three locations, corresponding to the generative process underlying a given sequence. In order to monitor the dynamics of inference, we used a continuous finger tracking system. Subjects also provided a detailed subjective report after each sequence (*Fig. 2b,c*). Our design therefore combines two aspects: detecting when a regularity emerges and identifying its type (i.e. statistic or rule).

Fig. 2 | Behavioural task. **a**, *Example sequences*. Sequences were of three different types: random from the beginning to the end or composed of an initial random part and a second non-random part that could either be produced by a statistical bias (here, a bias toward alternation) or a deterministic rule (here, the repetition of the AABB pattern). These sequences were composed of 200 low- and high-pitch tones (interval of 300 ms) that were randomly assigned to observations A and B at the beginning of each sequence. **b**, *Timecourse of a trial*. Each trial was self-initiated. During sequence presentation, subjects were asked to move their finger within a triangular arena whose vertices correspond to the three possible generative processes: random (H_{rand}), statistical bias (H_{stat}) or deterministic rule (H_{rule}). The mapping between left/right sides and statistics/rules was counterbalanced between subjects. Moreover, subjects answered several offline questions, notably about the moment when a regularity was detected. **c**, *Triangular arena*. The triangular arena provides a coordinate system where each location corresponds to a triplet of posterior probabilities of the three different hypotheses.
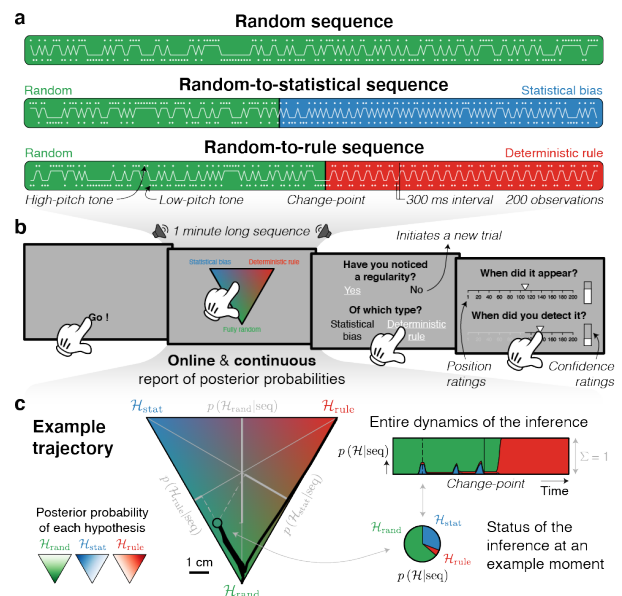
Normative two-system model. We formalized our theoretical proposal as an ideal observer model of the task. This normative model combines the sequence of observations with prior knowledge

about the task structure in order to constantly estimate, through Bayesian inference, the posterior

probability of the three possible generative processes (i.e. random, random-to-statistical and

random-to-rule). Under the random hypothesis, the model considers that observations are drawn

randomly without bias. For the non-random hypotheses, the model considers that there is a single

change-point that separates an initial random part from a second non-random part. For the random-

to-statistical hypothesis, the model considers that the second part is characterised by a bias in the

transition probabilities between successive observations. For the random-to-rule hypothesis, the

observer considers that the second part corresponds to the repetition of a particular pattern. The

position of the change-point, the bias in transition probabilities, and the repeated pattern are

unknowns that must be inferred from the sequence itself. Bayesian inference is used in all cases,

providing a unified account of regularity detection independently of the type of regularity. We

simulated this model with the same sequences that were presented to subjects in order to derive

quantitative predictions, and compare them to subjects' behaviour.

Offline, post-sequence, reports. *Generative process*. At the end of each sequence, in order to assess

subjects' sensitivity in detecting either type of regularity, we asked them to retrospectively judge

whether or not a regularity was present (detection) and, if so, to report its type (discrimination).

Subjects were able to both detect the presence of a regularity and to discriminate its type (*Fig. 3a*).

They performed above-chance in identifying each of the three generative processes (on average

61.7%, 71.9%, and 90.4%; chance = 33.3%, $d_{\text{Cohen}} > 1.21$, $t_{22} > 5.80$, $p < 7.76 \cdot 10^{-6}$). However,

accuracy differed across sequence types ($w^2 = 0.022$, $F_{2,22} = 20.6$, $p = 4.86 \cdot 10^{-7}$): subjects were

better at detecting deterministic rules than statistical biases (difference in accuracy = 18.5%,

CI = [11.4, 25.6], $d_{\text{Cohen}} = 1.12$, $t_{22} = 5.38$, $p = 2.12 \cdot 10^{-5}$).

To confirm that these results were not due to response biases, we relied on a (2D) signal

detection theory analysis and quantified subjects' sensitivities (*Fig. 3b*). As suggested by the

previous analysis, the detection of a statistical bias was accompanied by a lower sensitivity than the

detection of a deterministic rule (difference in $d' = 0.76$, CI $= [0.55, 0.96]$, $d_{Cohen} = 1.59$, $t_{22} = 7.62$, $p = 1.32 \cdot 10^{-7}$). We applied the same signal detection theory analysis to the model, and the generative hypothesis that it estimates as most likely at the end of each sequence. As in subjects, the model's sensitivity was lower for statistical biases compared to deterministic rules (difference in $d' = 0.39$, CI $= [0.20, 0.59]$, $d_{Cohen} = 0.89$, $t_{22} = 4.25$, $p = 3.31 \cdot 10^{-4}$).
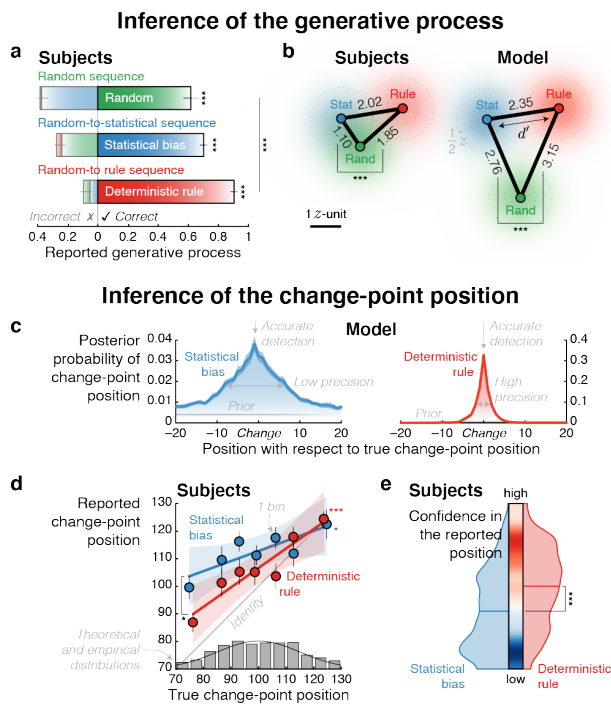
Fig. 3 | Offline (post-sequence) identification of the generative process and change-point position. a, *Subjects' post-sequence reports of sequence type.* Reported generative process is plotted as a function of the true generative process. **b**, *Signal detection theory analysis of sequence identification.* Multidimensional signal detection theory is applied on subjects reports in order to compare sensitivity to different regularities independently of possibly existing response biases. Sensitivity measures (*d'*, in *z*-units) are represented as edges of a triangle whose vertices represent the three possible generative processes. Similar to the model, the subjects are worse at distinguishing statistical biases from a random process than they are for the deterministic rules. **c**, *Model beliefs about change-point position.* The posterior distribution over change-point position (here, centred on true change-point position and averaged) reveals that the model accurately infers the position of the change-point but also that it is more precise at detecting the onset of a deterministic rule than it is at detecting the onset of a statistical bias. **d**, *Correlation between subjects' estimates of change-point positions and true positions.* Post-sequence subjects' estimates of change-point position as a function of true change-point position (in 7 equally spaced bins). Correlation is stronger in the case of deterministic rules compared to statistical biases. The shaded area corresponds to the 95% confidence interval of the regression coefficients. **e**, *Subjects' confidence in their estimates of change-point position.* Subjects use the lower part of the confidence scale (presented in *Fig. 2b*) more often in the case of statistical biases compared to deterministic rules, and the upper part of the confidence scale more often for deterministic rules than for statistical biases, leading to differences in average confidence. Distributions reflect kernel densities with bandwidth of 5%. In **c** to **e**, the analyses are restricted to non-random sequences that were correctly identified by subjects. In **a**, **d**, and **c**, error bars correspond to the standard error of the mean computed over subjects. Stars denote significance: *** $p < 0.005$, ** $p < 0.01$, * $p < 0.05$.

Interestingly, statistical biases that were missed by the subjects were not characterised by a late occurrence of the change-point (non-significant difference in change-point position $= -2.84$, CI $= [-9.03, 3.35]$, $d_{Cohen} = -0.20$, $t_{22} = -0.95$, $p = 0.35$, $BF_{null} = 3.05$) but instead by a lower posterior probability (i.e. closer to chance level) estimated by the model, relative to regularities that were detected (difference in $p(H_{stat}|seq) = 0.15$, CI $= [0.06, 0.25]$, $d_{Cohen} = 0.68$, $t_{22} = 3.27$,

p = 0.004). Thus, subjects missed regularities when the evidence was less strong. Although subjects were overall less sensitive than the model, those results indicate that the offline detection and identification of regularities is well accounted for by the model. We restrict the forthcoming analyses to regularities that were accurately detected in order to ensure that the effects we report are not due to a miscategorisation of sequences (*Table 1*).

| Condition	Type of regularity	Length	$p(A|B)$	$p(B|A)$	$p(A)$	$p(alt.)$	H	Subjects DR	Model DR
Random	— (x 10) [0]	—	1/2	1/2	1/2	1/2	2.00	14.2/23**	21.9/23***
Statistical bias	Repetition (low) [1]	—	1/3	1/3	1/2	1/3	1.92	16/23***	18/23***
Statistical bias	Repetition (medium) [2]	—	1/4	1/4	1/2	1/4	1.81	20/23***	23/23***
Statistical bias	Repetition (high) [3]	—	1/5	1/5	1/2	1/5	1.72	19/23***	23/23***
Statistical bias	Alternation (low) [4]	—	2/3	2/3	1/2	2/3	1.92	13/23*	19/23***
Statistical bias	Alternation (medium) [5]	—	3/4	3/4	1/2	3/4	1.81	15/23***	23/23***
Statistical bias	Alternation (high) [6]	—	4/5	4/5	1/2	4/5	1.72	14/23**	23/23***
Statistical bias	Frequency (low)	—	2/3	1/3	2/3	4/9	1.84	14/23**	17/23***
Statistical bias	Frequency (medium)	—	3/4	1/4	3/4	3/8	1.62	21/23***	23/23***
Statistical bias	Frequency (high)	—	4/5	1/5	4/5	8/25	1.44	19/23***	23/23***
Statistical bias	Frequency & repetition	—	1/2	1/4	2/3	1/3	1.79	20/23***	21/23***
Statistical bias	Frequency & alternation [7]	—	3/4	1/2	3/5	3/5	1.90	11/23(ns)	13/23*
Deterministic rule	AABB [0]	4	1/2	1/2	1/2	1/2	2.00	23/23***	23/23***
Deterministic rule	AAABBB [1]	6	1/3	1/3	1/2	1/3	1.92	23/23***	23/23***
Deterministic rule	AABABB [4]	6	2/3	2/3	1/2	2/3	1.92	22/23***	23/23***
Deterministic rule	AAABAB	6	1	1/2	2/3	2/3	1.59	21/23***	23/23***
Deterministic rule	AAAABBBB [2]	8	1/4	1/4	1/2	1/4	1.81	23/23***	23/23***
Deterministic rule	AABABABB [5]	8	3/4	3/4	1/2	3/4	1.81	23/23***	23/23***
Deterministic rule	AABBBABB [0]	8	1/2	1/2	1/2	1/2	2.00	20/23***	23/23***
Deterministic rule	AAAAABBBBB [3]	10	1/5	1/5	1/2	1/5	1.72	22/23***	23/23***
Deterministic rule	AABABABABB [6]	10	4/5	4/5	1/2	4/5	1.72	21/23***	23/23***
Deterministic rule	AAABAABBAB [7]	10	3/4	1/2	3/5	3/5	1.90	10/23(ns)	23/23***

Table 1 | Experimental conditions and detection rates. Superscripts indicate matched conditions. *H* corresponds to the Shannon entropy of transition probabilities (in bits). Detection rates (DR) denote the number of subjects who reported the correct generative process. The model identifies the generative process based on the maximum a posteriori probability over hypotheses at the end of the sequence. Stars denote significance of an exact binomial two-tailed test against ⅓: *** $p < 0.005$, ** $p < 0.01$, * $p < 0.05$; ns stands for non-significant.

Change-point position. Because of its hierarchical and rational nature, the model entertains a posterior distribution over possible change-point positions (*Supplementary Fig. 1*). Its maximum a posteriori estimate was highly correlated with the true change-point position across sequences (*Fig. 3c*), but less so in the case of statistical biases than for deterministic rules (difference in correlation = 0.53, CI = [0.38, 0.69], d_{Cohen} = 1.47, t_{22} = 7.05, p = 4.46 · 10^{-7}), indicating that, even for a normative model, the detection of a change-point is more accurate for an emerging deterministic rule than for an emerging statistical bias. In order to test those predictions, at the end of each sequence for which they reported the presence of regularity, we asked subjects to report the most likely position of the change-point, on a scale spanning the sequence length. As for the model, their estimates of change-point positions were correlated with the true positions for

statistics (correlation = 0.21, CI = [0.04, 0.38], d_{Cohen} = 0.53, t_{22} = 2.55, p = 0.018) and rules (correlation = 0.43, CI = [0.28, 0.57], d_{Cohen} = 1.24, t_{22} = 5.96, p = 5.34 · 10^{-6}), but significantly less for statistics (difference in correlation = 0.22, CI = [0.02, 0.41], d_{Cohen} = 0.49, t_{22} = 2.33, p = 0.029; *Fig. 3d*). In order to probe whether subjects were aware that their estimates of the change-point position were less accurate for statistical biases than for deterministic rules, we asked them to report their confidence in their estimates. Reported confidence was indeed lower for statistics compared to rules (difference in confidence = 14.4%, CI = [11.0, 17.8], d_{Cohen} = 0.53, t_{22} = 8.76, p = 1.26 · 10^{-8}). In the model too, the confidence about the change-point location, which can be formalized as the log-precision of the posterior distribution [19,44,45] over change-point positions, is also markedly lower for statistical regularities (difference in log-precision = 5.02, CI = [4.83, 5.20], d_{Cohen} = 11.7, t_{22} = 55.9, p = 3.30 · 10^{-25}; *Fig. 3e*).

Online report with finger tracking. We now turn to finger tracking recordings in order to explore whether subjects can faithfully track, in real time, the dynamics of their inference. Throughout the sequence, subjects continuously reported their current beliefs about the generative process of the observed sequence by moving their finger on a touch screen, within a triangular arena whose vertices correspond to each possible generative hypothesis: random, random-to-statistical, and random-to-rule (*Fig. 2c*). Importantly, each position in the triangle corresponds to a unique given set of posterior probability in the 3 possible hypotheses H_{rand}, H_{stat}, and H_{rule} (e.g. the bottom corresponds to (1, 0, 0), which is also the starting point, the centre to (⅓, ⅓, ⅓), etc.). Subjects' finger trajectories are thus directly converted into a timeseries of posterior probabilities ascribed to each hypothesis, which can be compared to the timeseries of posterior probabilities estimated by the *normative two-system model* (*Fig. 4a,b*).
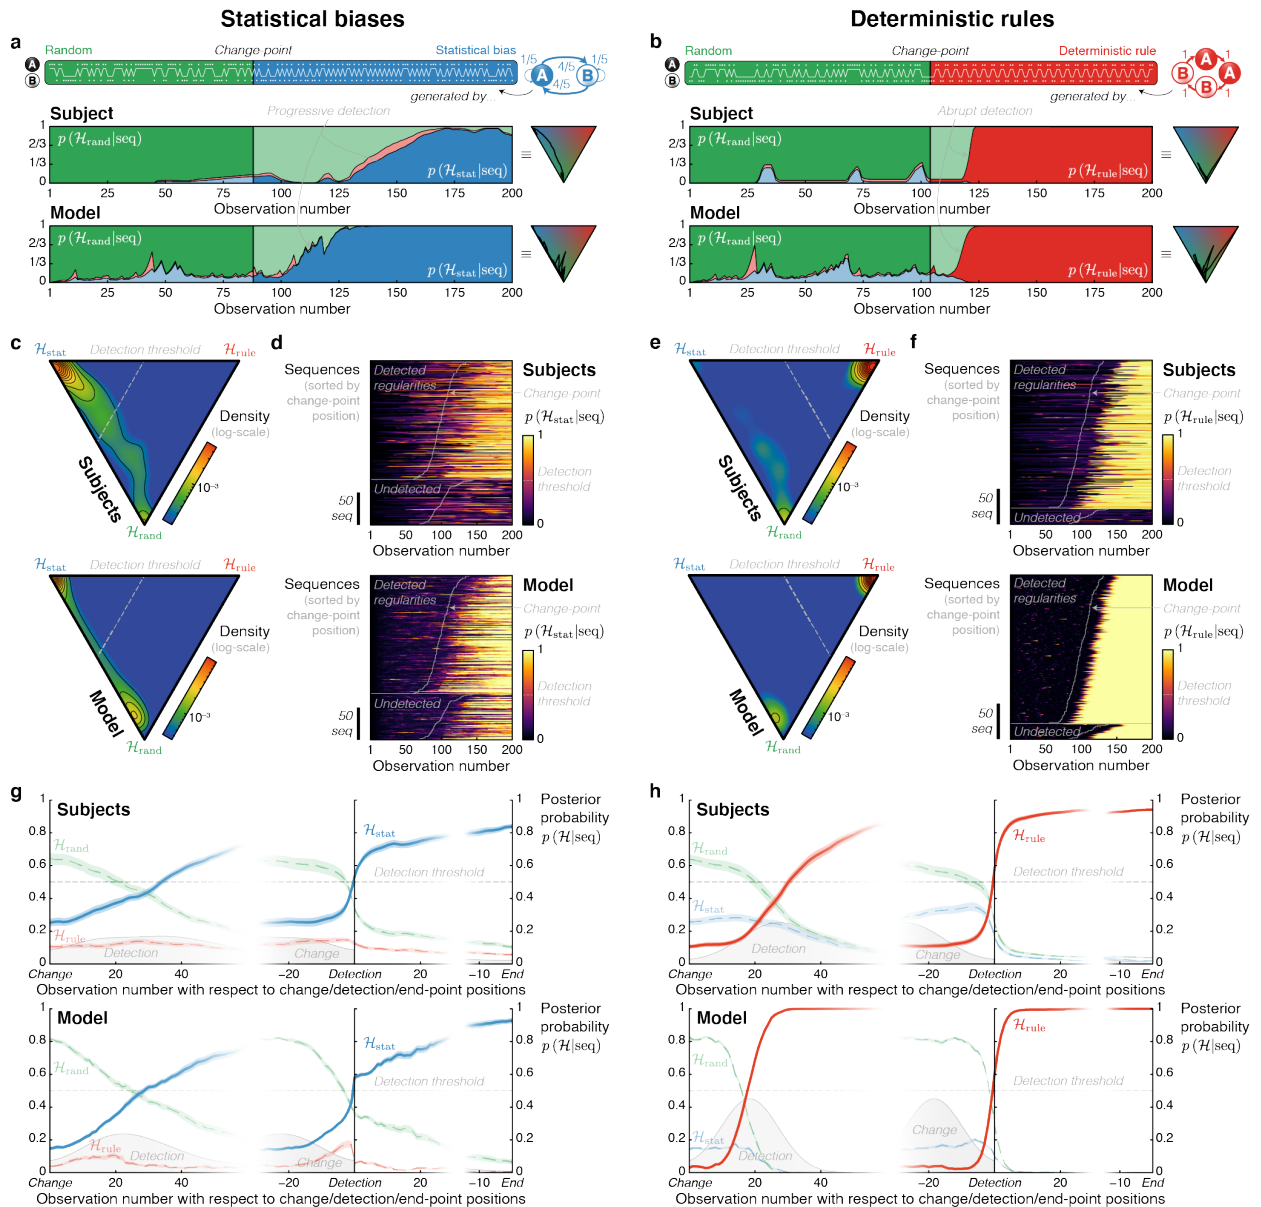
Fig. 4 | Different detection dynamics for statistical biases and deterministic rules. **a**, **b**, *Single-trial dynamics of hypothesis posterior probability* from an example subject and the model presented with the same random-to-statistical/rule sequence. **c**, **e**, *Triangular histograms* (smoothed and log-scaled) from the subjects and the model during parts of sequences with a statistical bias/deterministic rule. Different parts of the triangular arena are used in the different types of non-random parts of sequences. **d**, **f**, *Single-trial dynamics of the posterior probability in the statistical bias/deterministic rule hypothesis.* Sequences are grouped by detection (from subjects' post-sequence reports) and sorted by change-point position. Note the progressive/abrupt detection of statistical bias/deterministic rule. **g**, **h**, *Averaged dynamics of posterior probability locked on change, detection and end points of parts of sequences with a statistical bias/deterministic rule.* The average increase in the probability of the statistical bias/deterministic rule hypothesis also shows the progressive/abrupt difference, in subjects and the model alike. This difference is even clearer when the curves are time-locked to the detection-point defined as the moment when $p(H_{correct}|seq) > \frac{1}{2}$. Here, analyses were further restricted to non-random sequences that were correctly identified by subjects and for which detection-points were found for both the subjects and the model. Distributions reflect kernel densities of change/detection-points with a bandwidth of 8 observations. The shaded area corresponds to the standard error of the mean computed over subjects.

Dynamics of regularity detection. Under the hypothesis that subjects approximate the *normative two-system model*, we expect the dynamics of regularity detection to be different in the case of

statistical biases vs. deterministic rules. Below, we list those normative differences, illustrate them with the model and test them in subjects' finger trajectories.

Overall distribution of beliefs. We created 2D triangular histograms of finger position across all sequences in order to quantify belief distributions. Subjects, like the model, used the bottom part of the triangle, corresponding to the random hypothesis, as well as the parts of the triangle close to the statistical bias hypothesis vertex during random-to-statistical sequences and, similarly, the parts of the triangle close to the deterministic rule hypothesis vertex in the random-to-rule sequences (*Fig. 4c,e*); thereby demonstrating a correct identification of the generative process in the course of the sequence.

Build-up of beliefs. In order to explore the differences in belief update for statistical biases and deterministic rules, we now turn to the analysis of time-resolved trajectories. From now on, we further restrict analyses of non-random sequences to those for which we were able to find a detection-point (i.e. when the correct hypothesis becomes more likely than any others). In the model, the posterior probability of the correct hypothesis increased much more steeply in the case of rules than for statistics, resulting in step-like vs. gradual trajectories respectively. Subjects behave similarly, as can be seen on the trajectories locked on the true change-point position. However, locking on the true change-point underestimates the step-like nature of belief update in the case of deterministic rules due to variability in the onset of the step with respect to the change-point. We therefore also locked the trajectories onto the moment when the probability of the correct hypothesis crosses the detection threshold, which we refer to as the detection-point. This analysis revealed that in subjects, just as in the model, the increase in the posterior probability of the deterministic rule hypothesis is quite abrupt whereas the increase in the posterior probability of the statistical bias hypothesis is quite gradual. This difference is readily seen when inspecting individual trials (*Fig. 4d,f*). We quantified this difference by fitting sigmoid functions to trajectories in individual

trials and averaged the fitted slope parameter across sequences (*Supplementary Fig. 2*): it was significantly higher for rules compared to statistics in the model (difference in slope = 0.39, CI = [0.32, 0.46], $d_{Cohen} = 2.56$, $t_{22} = 12.3$, $p = 2.49 \cdot 10^{-11}$) and in subjects (difference in slope = 0.28, CI = [0.19, 0.37], $d_{Cohen} = 1.37$, $t_{22} = 6.55$, $p = 1.38 \cdot 10^{-6}$).

Regularity-specific modulations of detection dynamics. The *normative two-system model* predicts specific differences in detection dynamics within each type of regularity. More specifically, the strength of the statistical bias should modulate the amount of belief update, thus changing the slope of the gradual detection dynamics characteristic of statistical biases (*Fig. 5a*). By contrast, the length of the deterministic rule should change the moment of the detection-point, while maintaining the abrupt detection dynamics typical of deterministic rules (*Fig. 5b*).

Statistical biases. We sorted statistical biases according to the type and amount of bias they induce: one item could be more frequent than the other, or repetition could be more (or less) frequent than alternation, or both of these two biases could be present. We quantified the strength of this bias by the Shannon entropy of the corresponding generative transition probabilities, which ranged from 1.44 to 1.92 bit (*Fig. 5c*). All statistical biases were detected gradually, but the amount of belief update varied across regularities (*Fig. 5d*): with a significant effect of entropy for the model (correlation = –0.28, CI = [–0.41, 0.14], $d_{Cohen} = –0.90$, $t_{22} = –4.33$, $p = 0.0003$) and the subjects (correlation = –0.32, CI = [–0.49, 0.15], $d_{Cohen} = –0.82$, $t_{22} = –3.91$, $p = 0.0008$).

Deterministic rules. We distinguished among deterministic rules according to their lengths (4, 6, 8 or 10 observations; *Fig. 5e*). Furthermore, some of the patterns purposely comprised a bias in the apparent transition probabilities between items (*Fig. 5g*). For instance, the repetition of the AAAABBBB pattern results in many more repetitions (75%) than expected by a random process (50%). By contrast, the repetition of AABB produces as many As and Bs (50%) and as many

repetitions as alternations (50%), which is equal to what is expected on average from a random process. The results showed that the detection-point (*Fig. 5f*) increased with the length of the pattern for both the model (correlation = 0.87, CI = [0.79, 0.95], d_{Cohen} = 4.48, t_{22} = 21.5, $p = 2.99 \cdot 10^{-16}$) and the subjects (correlation = 0.51, CI = [0.43, 0.59], d_{Cohen} = 2.76, t_{22} = 13.3, $p = 5.80 \cdot 10^{-12}$). Importantly, even though the detection-point varied according to the length of the patterns, the detection dynamics always remained similarly abrupt (*Fig. 5h*); in particular belief update was not modulated by the strength of the statistical bias induced by the pattern (quantified by Shannon entropy), neither in the model (non-significant correlation = 0.10, CI = [–0.068, 0.27], d_{Cohen} = 0.26, t_{22} = 1.23, p = 0.23, BF_{null} = 2.35) nor in subjects (non-significant correlation = 0.078, CI = [–0.071, 0.23], d_{Cohen} = 0.23, t_{22} = 1.08, p = 0.29, BF_{null} = 2.72).

Fig. 5 | Regularity-specific rational modulation of detection dynamics. **a**, *Examples belief updates from the model*. The amount of belief update, that is the difference in belief between the change-point and the end of the sequence, is expected to vary depending on the statistical bias considered. **b**, *Example detection-points from the model*. The amount of belief update is not expected to vary depending on the deterministic rule considered, only the detection-point, that is the number of observations between change- and detection-points (i.e. when the deterministic rule hypothesis is more likely than any other hypothesis). **c**, *Types of statistical biases*. Statistical biases are defined by first order transition probabilities, which have different subtypes of bias (item frequency, repetition-alternation frequency) and strength (measured by Shannon entropy). **d**, *Modulation of belief update for statistical biases*. Change-points, if detected, should lead to belief update. In random-to-statistical sequences, the amount of belief update decreased with Shannon entropy of generative first order transition probabilities. **e**, *Types of deterministic rules*. Deterministic rules are defined by the pattern that is repeated, which has a specific length, and which induces a specific statistical bias. **f**, *Modulation of detection-point for deterministic rules*. Change-points, if detected, should led to crossing of the detection threshold. In random-to-rule sequences, the detection-point increased with pattern length. **g**, *Apparent statistical biases of deterministic rules*. Repeating patterns induce, on purpose, different types of statistical biases, of different strength (measured as Shannon entropy). Most of them are matched with biases induced by pure statistical biases. **h**, *No modulation of belief update for deterministic rules.* Contrary to pure statistical biases, belief update of deterministic rules does not decrease with Shannon entropy of apparent first order transition probabilities. In **d**, **f** and **h**, each dot is a regularity; and the analyses were restricted to regularities correctly classified by subjects and for which detection-points were found for both the subjects and the model. Error bars correspond to the standard error of the mean computed over subjects. Stars denote significance: *** $p < 0.005$, ** $p < 0.01$, * $p < 0.05$.
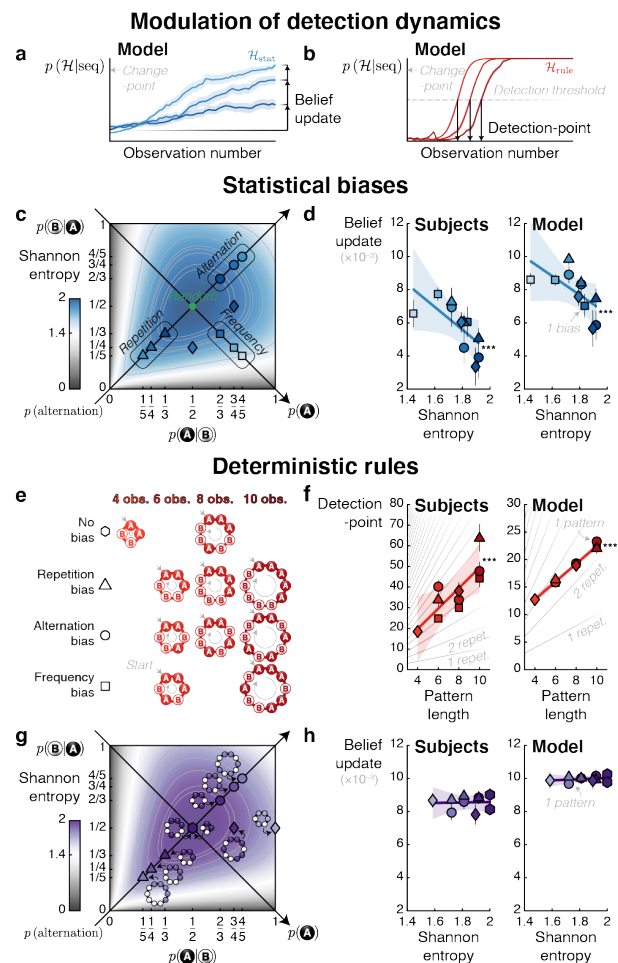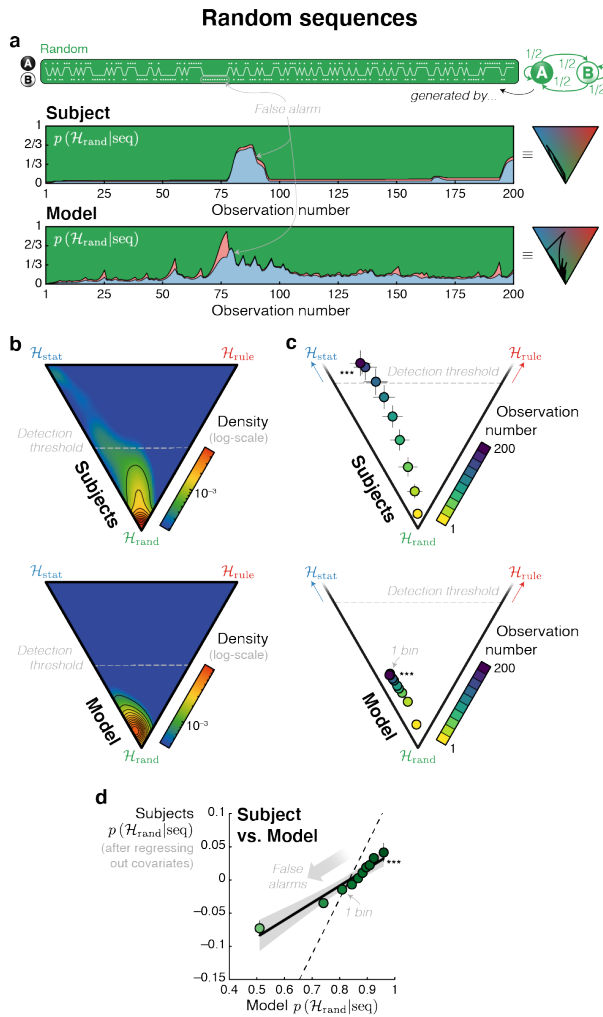
False alarms in random sequences. Throughout random sequences, the *normative two-system model* predicts that observers keep track of all hypotheses. Thus, they should entertain high probabilities in the random hypothesis, as indeed observed, but because random contains transient periods of apparent regularity, their beliefs in non-random hypotheses should also occasionally increase (*Fig. 6a*). The results indicate that such false alarms do occur in subjects. Furthermore, in both the model and the subjects, they much more often concerned a transient preference for the statistical bias hypothesis than for the deterministic rule hypothesis (*Fig. 6b*). This is because local trends, even weak, constantly impact the statistical hypothesis by slightly shifting the estimated probabilities away from chance. By contrast, transient trends impact the deterministic hypothesis only when they license strong (erroneous) predictions, which rarely occur by chance. Below, we explore the dynamics and rationality of subjects' false alarms during random sequences (i.e. the initial part of random-to-non-random sequences and the entire length of random sequences).

We first tested whether subjects' false alarms increase with the number of observations received (*Fig. 6c*). This is a prediction of the model: the posterior probability of the random hypothesis steadily decreases with the number of observations in the sequence (correlation = –0.32, CI = [–0.34, –0.29], d_{Cohen} = –5.53, t_{22} = –26.5, p = 3.45 · 10^{-18}). This is a cornerstone of hierarchical inference in our task: the model assumes that a change-point occurs in about two thirds of sequences, which translates into a slowly increasing probability that a change-point has occurred as the sequence unfolds in time. Subjects show the same effect (correlation = –0.60, CI = [–0.67, –0.54], d_{Cohen} = –4.18, t_{22} = –20.1, p = 1.24 · 10^{-15}).

We then tested whether false alarms occurred precisely when the model predicted them, that is when the random sequences exhibited, by chance, a transient regularity. To do so, posterior probabilities in the random hypothesis reported by subjects were regressed against the model's posterior probabilities in the random hypothesis (*Fig. 6d*). We found a significant positive effect (coefficient = 0.58, CI = [0.50, 0.65], d_{Cohen} = 3.34, t_{22} = 16.0, p = 1.27 · 10^{-16}) that also survived when the number of observations received was included as covariate (coefficient = 0.23,

CI = [0.16, 0.30], d_{Cohen} = 1.41, t_{22} = 6.77, p = 8.41 · 10^{-7}). This result confirms that subjects' false

alarms reflect a rational inference due to transient deviations from randomness, over and above the

simple effect of elapsed time since the beginning of the sequence.

Fig. 6 | False alarms reflect transient periods of regularity. a, *Example random sequence and resulting inferences.* Transient periods of regularity can appear, by chance, in random sequences (i.e. the initial part of random-to-non-random sequences, or the entire length of the random sequences). They can lead observers to make false alarms, that is, decrease the posterior probability of the random hypothesis in favour of a non-random hypothesis. **b**, *Triangular histograms* (smoothed and log-scaled) for the subjects (top panel) and the (bottom panel) during random sequences. **c**, *Effect of observation position within the sequence on false alarms.* Reported probabilities from the subjects and the model are averaged in groups of 20 consecutive observations in the sequence. The false alarm rate increased as more observations were received within a sequence, in subjects (top panel) and the model (bottom panel). **d**, *Correlation between subjects' and model's posterior probability in the random hypothesis.* A linear relationship between subjects and the model's probability of the random hypothesis (in 10 bins defined using deciles) remains even after regressing-out the confounding effect of time elapsed in the sequence from subjects' data (dashed line: without regressing-out). The shaded area corresponds to the 95% confidence interval of the regression coefficients. In **c** to **f**, error bars correspond to the standard error of the mean computed over subjects. Stars denote significance: *** $p < 0.005$, ** $p < 0.01$, * $p < 0.05$.

Rational comparison between non-random hypotheses. A crucial assumption of the *normative*

two-system model is that subjects perform a Bayesian inference over two distinct hypothesis spaces

in order to detect and identify statistical biases and deterministic rules. The normative nature of this

inference enables to compute the posterior probability of the different non-random hypotheses in

the same probabilistic currency and, therefore, to compare them. To study this prediction, we now

explore situations of conflict, when both non-random hypotheses compete without a clear winner.

Those situations appear frequently just before the detection of a deterministic rule as well as during random sequences. We used the belief difference between non-random hypotheses (the difference between their posterior probability normalized by their sum) as an index of conflict between statistics and rules in those periods.

Before the detection of deterministic rules, the model often shows a period of indecision during which non-random hypotheses have intermediate posterior probabilities. Those intermediate levels temporarily favour the statistical bias hypothesis before a change-of-mind [46] leads to an accurate identification of the deterministic rule, and does so all the more that the repeated pattern has a strong bias in its apparent transition probabilities (*Fig. 7a*). To assess this effect, we correlated the strength of this bias (quantified by Shannon entropy) with the belief difference averaged during the period going from the true change-point to the detection-point in the model (correlation = 0.50, CI = [0.38, 0.62], d_{Cohen} = 1.81, t_{22} = 8.69, p = 1.46 · 10^{-8}) and in the subjects (correlation = 0.19, CI = [0.0035, 0.38], d_{Cohen} = 0.44, t_{22} = 2.11, p = 0.046). Importantly, the belief difference assessed in the model and in subjects' data during this indecision period correlated with each other across patterns (correlation = 0.22, CI = [0.070, 0.37], d_{Cohen} = 3.04, t_{22} = 3.04, p = 0.006; *Fig. 7b*).

During random sequences, while the model generally assigns a low probability to both statistics and rule hypotheses, their relative posterior probability fluctuates depending on the exact observations received. We reasoned that if subjects used a rational arbitration between non-random hypotheses, the relative credence they assign to both hypotheses should show fluctuations similar to the model. We therefore regressed subjects' difference in beliefs between non-random hypotheses against the model's difference in beliefs (*Fig. 7c*). We found a significant positive effect (coefficient = 0.28, CI = [0.21, 0.35], d_{Cohen} = 1.71, t_{22} = 8.18, p = 4.08 · 10^{-8}) that also survived when confounding variables were included in the regression model: the posterior probability of the random hypothesis alone (coefficient for belief difference = 0.24, CI = [0.18, 0.31], d_{Cohen} = 1.55, t_{22} = 7.41,

$p = 2.04 \cdot 10^{-7}$) or together with its interaction with the belief difference

(coefficient for belief difference $= 0.17$, CI $= [0.02, 0.33]$, $d_{\text{Cohen}} = 0.49$, $t_{22} = 2.33$, $p = 0.029$).
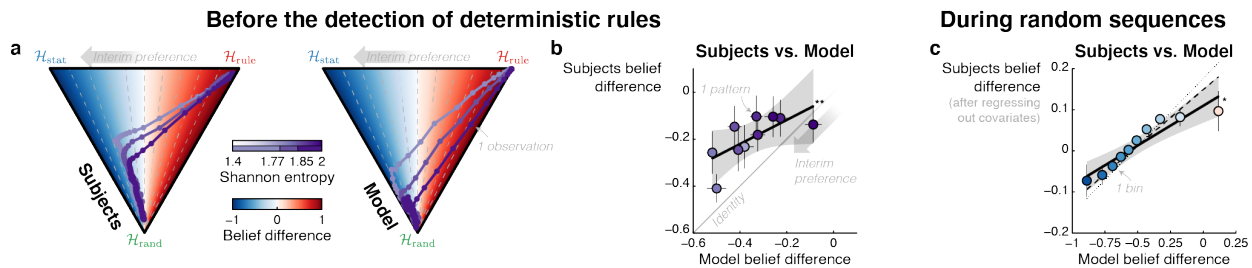
Fig. 7 | Rational competition between non-random hypotheses. **a**, *Interim preference for the statistical bias hypothesis before the detection of deterministic rules depends upon the entropy of the apparent transition probabilities characterising the pattern.* Trajectories during random-to-rule sequences were centred on the detection-point and binned according to the entropy of the repeated pattern. Each dot corresponds to one observation in the sequence. The belief difference is the difference between the posterior probability of the two non-random hypotheses normalized by their sum (it evolves between –1 and 1 independently of the probability of the random hypothesis), which is shown here as a heatmap. **b**, *Belief difference from subjects as a function of the belief difference from the model across patterns*. The belief difference between non-random hypotheses was averaged from the change-point to the detection-point, in both subjects and the model. Each dot corresponds to one pattern, whose entropy is color-coded. In **a** and **b**, analyses were restricted to random-to-rule sequences that were correctly classified by subjects and for which detection-points were found. **c**, *Belief difference from subjects as a function of the belief difference from the model in random sequences.* In random sequences (i.e. the initial part of random-to-non-random sequences, or the entire length of the random sequences), the belief difference from subjects was regressed against the model belief difference. The linear relationship between subjects and the model's belief difference is shown (in 10 bins defined using deciles) after regressing-out effect of confounding variables such as the posterior probability of the random hypothesis (dotted line: without regressing-out), and together with its interaction with the posterior difference in the regression (dashed line: without regressing-out). In **b**, and **c**, error bars correspond to the standard error of the mean computed over subjects and the shaded area to 95% confidence interval of the regression coefficients. Stars denote significance: *** $p < 0.005$, ** $p < 0.01$, * $p < 0.05$.

Alternative models. The *normative two-system model* accounts for many important aspects of the

human detection and identification of temporal regularities observed in the current experiment. To

test the necessity of the model's assumptions, we now explore alternative models lacking one or the

other of its fundamental properties (*Fig. 8a*). We test, first, a model with a single common system

for deterministic rules and statistical biases (instead of two hypothesis spaces); and, second, a model

with distinct but non-commensurable systems (instead of the shared probabilistic currency).

Normative single-system model. A first alternative to the *normative two-system model* is that

subjects use a continuum rather than two distinct hypothesis spaces (hence *single-system* instead of

two-system): because deterministic rules are a limit case of statistical biases (with $p = 0$ or 1),

subjects could employ statistical learning in all cases.

For completeness, we tested several versions of this *normative single-system model.* A first version monitors order 1 transition probabilities and distinguishes between statistics vs. rules based on the strength of the inferred probability bias (intermediate vs. strong, respectively). However, this model has obvious peculiarities that discard it; for instance, unlike subjects (*Table 1*), it cannot detect the AABB pattern, whose repetition causes no bias in the apparent first-order transition probabilities.

A second version of the *normative single-system model* monitors transition probabilities of order 1 vs. higher-order (up to 9 corresponding to the longest patterns used) for statistical biases vs. deterministic rules respectively. This is plausible a priori because the repetition of a pattern necessarily induces a bias in the apparent higher-order transition probabilities (e.g. for the pattern AABB, the order 2 transition probabilities are $p(B|AA) = p(A|BA) = p(B|AB) = p(A|BB) = 1$). For every order considered, this model predicted that the detection of deterministic rules should be characterized by gradual rather than abrupt detection dynamics (*Fig. 8b*), unlike subjects (this effect is explained in detail in *Supplementary Note 1*). This model thus provided a worse account of human data compared to the *normative two-system model* (difference in MSE $> 2.21 \cdot 10^{-5}$, $d_{\text{Cohen}} > 0.88$, $t_{22} > 4.22$, $p < 3.55 \cdot 10^{-4}$; *Fig. 8c*). In addition, low-order versions simply missed patterns that induce weak (if any) biases in transition probabilities of the corresponding order; while higher-order versions showed too many false alarms in favour of the deterministic rule hypothesis.

We also tried to rescue the *normative single-system model* that monitors transition probabilities of different orders using a prior distribution biased towards extreme values of 0 or 1 specifically for higher-order transition probabilities (H'_{rule}) [47], so as to capture the fact that deterministic rules afford more certain predictions than statistical biases. However, this version further aggravated the problem of excessive false alarms even for very weak biases of prior distributions (see *Supplementary Note 2*) and it could not mimic the observed abrupt detection of deterministic rules.

Non-commensurable two-system model. The *normative two-system model* considers the three hypotheses together (i.e. H_{rule}, H_{stat}, and H_{rand}) and assigns to each of them a posterior probability that takes into account the other two hypotheses (through normalization in Bayes' rule). A second alternative to this *normative two-system model* is that subjects use distinct hypothesis spaces (i.e. H_{rule}, H_{stat}, and H_{rand} as before) but estimate the posterior probabilities of each non-random hypothesis only relatively to the random hypothesis, while ignoring the third hypothesis (i.e. H_{rule} vs. H_{rand} and H_{stat} vs. H_{rand}). In this case, owing to a lack of normalization across the three hypotheses, their probabilities do not sum to 1 and cannot be directly compared (hence *non-commensurable* instead of *normative*), we will thus term those probabilities pseudo posteriors. The *non-commensurable two-system model* must resort to a non-normative way of comparing hypotheses in order to report a belief in the triangular arena. We explored several ways.

A first version of this *non-commensurable two-system model* obeys a max rule: (1) it selects among the two putative regularities based on the maximum pseudo posteriors, and (2) it selects between the selected regularity and the random hypothesis according to what the corresponding pseudo posterior supports. The average detection dynamics of this model resembles those of the *normative two-system model* (*Fig. 8d*), but comparison with single-trial dynamics shows that the *normative two-system model* provides a better account of the subjects' trajectories (difference in MSE = 0.014, CI = [0.003, 0.025], d_{Cohen} = 0.53, t_{22} = 2.55, p = 0.018; "maximum" in *Fig. 8e*). This difference arises because the max rule generates sudden jumps between the left and right edges of the triangle and never occupies the centre, which is a landmark signature of the relative weighing between non-random hypotheses.
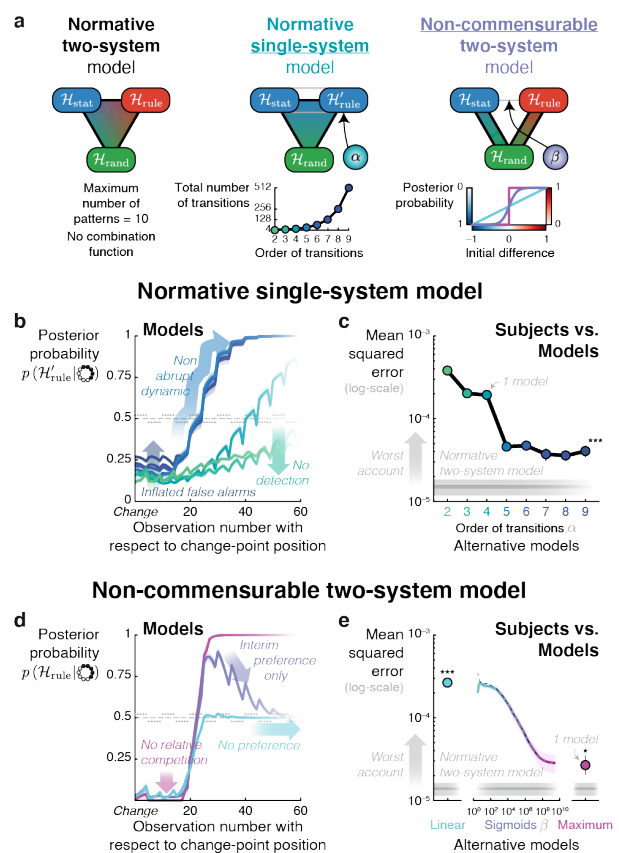
A second version of this *non-commensurable two-system model* weighs the two non-random hypotheses according to the (linear, i.e. exact) difference in their pseudo posterior probabilities, rather than selecting one in an all-or-none manner. The problem with this version is that it will often remain undecided, converging to a conclusion only if one of the two pseudo posterior probabilities vanishes (becomes 0), which does not happen whenever deterministic rules have a bias in their

apparent transition probabilities (e.g. AAAAABBBBB; see purple curve in *Fig. 8d*). Such

indecision does not occur in the *normative two-system model* because it correctly compares the two

non-random hypotheses based on their true posteriors (rather than pseudo posteriors). Again, the

non-commensurable two-system model resulted in a worse account of human data compared to the

normative two-system model (difference in MSE = 0.25, CI = [0.22, 0.28], d_{Cohen} = 3.59,

t_{22} = 17.23, p = 2.92 · 10^{-14}; "linear" in *Fig. 8e*).

Fig. 8 | Rejecting alternative models on the basis of deterministic detection dynamics. **a**, *Different models*. The normative two-system model is sketched side-by-side with two alternative models that either use a single common system for deterministic rules and statistical biases (instead of two distinct hypothesis spaces), or use distinct but non-commensurable systems (instead of the shared probabilistic currency). In the first case the α parameter defines the order of transition probabilities that are used to detect the deterministic patterns; it induces an exponential growth of the number of patterns to monitor. In the latter case, the β parameter controls how the non-commensurable systems are combined such that they can be reported in the triangular arena. **b**, *Detection dynamics of the AAAAABBBBB pattern by different versions of the normative single-system model*. Posterior probability in the deterministic rule hypothesis for different versions of the normative single-system model learning transition probabilities of different orders (from 2 to 9) is displayed after the change-point initiating the repetition of the AAAAABBBBB pattern. **c**, *Error of the different versions of the normative single-system model*. The mean squared error (MSE) between subjects' and models' abruptness (averaged over patterns) measures how well each model accounts for human data. A larger error indicates a worst account of human detection dynamics of deterministic rules. **d**, *Detection dynamics of the AAAAABBBBB pattern by versions of the non-commensurable two-system model*. Same as in **b** but in the case of different versions of the *non-commensurable two-system model* that combine the independently-computed posterior probabilities in the non-random hypotheses differently: a linear function of the difference, by selecting the best hypothesis and discarding the other, or a sigmoid function (here with an example using $\beta = 10$) of the difference or of the log-ratio. **e**, *Error of the different versions of the non-commensurable two-system model*. Same as in **c** but with the different versions of the *non-commensurable two-system model*. The dashed line corresponds to versions with a sigmoid-based combination but using the log-ratio between posterior probabilities of the non-random hypotheses (instead of the difference). In **a** to **e**, analyses were restricted to random-to-rule sequences that were correctly classified by subjects and for which detection-points were found. Error bars and shaded areas correspond to the standard error of the mean computed over subjects. Stars denote significance: *** $p < 0.005$, ** $p < 0.01$, * $p < 0.05$.

As an attempt to overcome the indecision problem, we tested a third version of the *non-*

commensurable two-system model that can exaggerate the difference in pseudo posterior

probabilities of the non-random hypotheses. We used a sigmoid function to formalize such

exaggeration, because it allows a significant impact of even subtle differences in pseudo posterior probabilities. Varying the slope of this sigmoid actually moves the model along a continuum whose ends correspond to the "max rule" and the "linear difference" versions tested above. However, this version suffered from the same problems as these previous versions (*Fig. 8d*) and no matter its parameterization, it always provided a worse account of subjects' behaviour compared to the *normative two-system model* (*Fig. 8e*). Note finally that resorting on a sigmoid that uses the log-ratio of (instead of the difference in) the pseudo posterior probabilities of the non-random hypotheses leads to the same conclusions (see dashed line in *Fig. 8e*).

Summary. The *normative two-system model* provided a better account of human behaviour than the many different alternative models we considered (additional arguments against these alternative models are presented in *Supplementary Note 2* and *Supplementary Fig. 4 to 7*), in spite of the fact that some of those alternative models had free parameters giving them more flexibility than the (parameter-free) *normative two-system model*.

Discussion

We proposed here that a simple taxonomy may help to organise the theoretically infinite number of temporal regularities that exist in the environment. This taxonomy distinguishes deterministic rules, for which certainty can be attained, from statistical biases, for which a degree of uncertainty remains. We asked human subjects to detect and identify regularities of either category that could suddenly appear within a random sequence of binary observations. We measured the dynamics of their inference by continuously tracking subjects' finger movement [46,48] and compared them to normative models [49,50]. We found that subjects assigned probabilistic credence to each of three distinct hypotheses, and that they treated statistics and rules as distinct hypothesis spaces rather than as a continuum, following our taxonomy.

Many previous studies focused on how participants use previous observations to make predictions about the future. Those predictions can leverage, often in different studies, either the estimation of a particular statistics (e.g. item frequency and transition probabilities) or the identification of a rule (e.g. a repeating pattern). We would like to stress that in the context of predictions the characterization of a specific generative process (estimation of a statistical bias or the identification of a rule) is a different computational problem than the identification of the most relevant process (rule or statistics ?) for predictions. The latter has received much less attention than the former, and it is the problem we tackled here.

A few previous studies approached this arbitration problem. For instance, some studies on the perception of randomness asked human subjects to categorize sequences as random versus non-random. Interestingly, such studies have shown that humans are prone to seeing regularity even in fully random sequences [51,52], which is similar to the false alarms we observed here during random sequences. We proposed that such false alarms may actually result from an optimal inference process. In another line of research, rule learning paradigms also prompted subjects to categorize random vs. non-random sequences [10]. In particular, a recent study [31] found that this categorisation relies on the computation of a posterior distribution over possible patterns (of different lengths), a finding we replicated here, while also extending it to other generative processes.

Another innovation of our work is to study the advent of regularity. Real-world environments are usually volatile, causing frequent changes in the process generating the observations. As a consequence, in addition to categorising among regularities, observers trying to predict future observations must also detect *when* regularity appears. We found that both statistics and rules could be promptly detected, but with different detection dynamics. Deterministic rules were quite abruptly detected, reflecting "aha moments" [53]; a computational consequence to the all-or-none predictions they afford. The onset of the detection varied as a function of pattern length. Both findings replicate previous studies in which repeating patterns appeared suddenly in otherwise random sequences [10,11]. By contrast, statistical biases were detected more progressively, reflecting

evidence accumulation [23,46]; a computational consequence of their inherent uncertainty. The slope of the detection varied as a function of the strength of the bias (quantified in terms of Shannon entropy). Again, both findings replicate previous studies [12,14].

Importantly, detecting the onset of a regularity, be it deterministic or statistical, requires observers to explicitly monitor change-points separating random from non-random observations. Because change-point positions are a priori unknown, they must be inferred from the sequence, and the uncertainty associated with this inference of position should be factored into the evaluation of the non-random hypotheses, which is a form of hierarchical inference. Knowing that change-point occurs in ⅔ of sequences here, the a priori probability of having encountered a change-point increases as the sequence unfolds in time. Accordingly, in random sequences, subjects' false alarms increased with sequence duration, as also observed in a recent study [12]. A model lacking such an explicit representation of change-points would, by contrast, have a stable false alarm rate throughout the sequence [54].

Subjects were also able to explicitly report change-point positions after sequence presentation, and rate their confidence in that estimation, both in a manner conforming to a hierarchical inference. This replicates the findings of recent studies on statistical learning with change points [15,17,19,20,55], but here we extended those findings to the case of deterministic rules.

Our results indicate that the human detection and categorisation of regularities rely on three different hypothesis spaces: a fully random hypothesis, a statistical bias hypothesis, and a deterministic rule hypothesis, which we formalized as the *normative two-system model*. We wish to dissipate a potential misunderstanding. By design, our task is to categorise sequences, using these three categories. The overall accurate categorisation that we observed is not diagnostic of an inference relying on different hypothesis spaces, since each category could correspond to different parameters within the same space (i.e. the same generative model of observations). We explored such a possibility with the *normative single-system model* in which statistics and rules both correspond to biases in transition probabilities, of low and high orders respectively. However, the

normative two-system and *normative single-system models* make different predictions regarding the pattern of categorisation errors, the rate of false alarms, and the dynamics of regularity detection. Crucially, a careful analysis of those aspects showed that subjects' behaviour conformed specifically to the *normative two-system model*: subjects treated statistics and rules as pertaining to fundamentally different hypothesis spaces.

The existence of distinct hypothesis spaces for statistics vs. rules seems to be supported by partly dissociable neural underpinnings. Studies on statistical learning and behavioural stochasticity typically report the involvement of a widespread network of brain regions including sensory cortices and the cingulate cortex [56,57], as well electrophysiological signatures with various (both early and late) latencies [18,58,59]. By contrast, sequences generated using artificial grammars or pattern repetitions and their violation typically recruit a more restricted set of brain regions including the inferior frontal gyrus and the hippocampus [10,30,33,36,56,60], and often elicit exclusively late electrophysiological signatures such as the P300 [58,59]. The few studies that have directly contrasted both types of regularities, yet not fully orthogonally, confirm a dissociation in terms of anatomy [22,32,37–39] and timing of evoked responses [58,59] associated with rules and statistics.

Using different hypothesis spaces for regularity detection, as opposed to using a continuum, has important computational advantages for biological agents. Firstly, it allows for a faster detection of deterministic rules. This is because the deterministic rule hypothesis considers only extreme probability values (essentially 0 vs. 1), such that evidence for or against it accumulates rapidly [23,46]. By contrast, the biases in higher-order transition probabilities that the *normative single-system model* tracks in order to detect rules take much longer to be detected.

Secondly, dividing the continuum of predictability into discrete hypotheses (as the *normative two-system model*) provides a tractable solution to the problem of regularity detection [61,62]. By contrast, computing with the continuum (as the *normative single-system model*) is either very difficult, if not simply intractable. One implementation of this solution, currently explored in the general artificial intelligence literature [63], is the mixture of local expert

agents [18,54,55,64]. This solution combines different algorithms (i.e. hypothesis spaces) that are specialized for different aspects of the inference process. When well-tuned, this solution can prove much faster and simpler than the full optimal solution [65,66].

Thirdly, such a division of tasks leverages dedicated computations for each hypothesis space, which can also reduce the total computational cost [67]. In our case, the distinction between statistics and rules enables both a simple, accurate estimation of low-order statistical biases *and* a sensitivity to long-distance dependencies, with pattern matching. Tracking higher-order statistics, as in the *normative single-system model*, also preserves a sensitivity to long-distance dependencies but at the cost of an exponential growth of the number of transitions to monitor with the length of the dependency. By contrast, the number of patterns that are monitored at a given moment by the *normative two-system model* remains small. This is because all patterns that are incompatible with the sequence are excluded, resulting in zero or one correct pattern per pattern length. Such a pruning strategy is very effective at reducing the cost of computations [25,61,62,68].

When using multiple hypothesis spaces, one must identify the one that is best suited for prediction or behavioural control, which requires arbitrating among them [69–71]. We used a behavioural apparatus and instructions that advantageously left subjects unconstrained in the way they would arbitrate between the different hypotheses and found that they weighed these hypotheses in a graded manner: (1) the random vs. statistical bias hypotheses graded comparison was demonstrated by the specific detection dynamics of statistical biases and the false alarms, (2) the random vs. deterministic rule hypotheses graded comparison by the specific detection dynamics of rules, (3) and the statistical biases vs. deterministic rule hypotheses graded comparison by an interim competition before the detection of rules and by the relative weighing between them during random sequences.

These aspects of behaviour were well accounted for by the *normative two-system model*, thereby showing that subjects use a common probabilistic currency to arbitrate between the different hypotheses, in line with previous accounts of Bayesian inference applied to discrete states [22,31,68].

We also explored alternative possibilities in the form of the *non-commensurable two-system model* which computes pseudo posterior probabilities of the non-random hypotheses independently of each other and, hence, cannot normatively compare them. We showed that the *normative* and *non-commensurable models* make different predictions regarding the detection dynamics of deterministic rules which rejected the latter. Worth noting, the only mathematical difference between the *normative two-system model* and the *non-commensurable two-system model* is the normalization term in Bayes' rule which either includes the three hypotheses, or only two of them, respectively.

This rational weighing of hypotheses is all the more useful that both types of regularities can coexist in the same input (*Fig. 1b*), including speech [72,73]. Deterministic rules can induce apparent statistical biases (e.g. the repetition of the AAB pattern induces globally more As than Bs), while statistical biases can induce local apparent rule-like regularities (e.g. …ABABAB… in frequently alternating sequences). Many previous studies, by investigating statistics and rules separately, have neglected the possibility that one type of learning could interfere with the other, or foster regularity detection. Here, by manipulating the extent to which deterministic rules also induce apparent statistical biases, we found an interim competition between statistics and rules, even causing changes-of-mind [46], and characterized it (with Shannon entropy).

We now acknowledge several limitations of our study. A first limitation is that we used a restricted set of regularities, due to experimental time constraints. This restricted set enabled us to study the estimation of first-order transition probabilities and detection of repeating patterns, but is not suited to study more complex types of regularity [2]. Notably, humans can detect deterministic rules in the form of algebraic patterns (e.g. AB-AABB-AAABBB… = $A^n B^n$ where n increases sequentially). A possibility, which has received some behavioural support, is that the human brain uses a "language of thought", including such algebraic patterns, in order to compress deterministic sequences into a regular expression [29,35,74–76]. For simplicity, we manipulated here only repeated patterns of various lengths, but found, in line with a possible role of compression, that the barely

compressible length-10 pattern AAABAABBAB was more often missed by subjects compared to the other two length-10 patterns, AABABABABB and AAAAABBBBB, which are more compressible ($[A^2, (BA)^2, B^2]$ and $[A^4, B^4]$ respectively). Note, however, that this effect could also result from the difference in the patterns' apparent statistical bias (which is respectively weak and strong), an alternative we have recently considered and discarded [76].

This brings us to the more general question of what is psychologically general in our model and what is specific to the task. Although we can only speculate about the scope of statistical biases (i.e. which order(s) of transitions) or deterministic rules (i.e. which type(s) of rules) that is relevant for the brain, our model makes the general claim that there exist fundamentally two categories of regularity, statistics and rules, which afford respectively uncertain vs. sure predictions. The model also posits that the statistical bias and deterministic rule hypotheses are evaluated in parallel during sequence processing, leading to testable predictions [76]. The model also posits that humans can rationally compare those different hypotheses using a common probabilistic currency. Future studies should thus investigate a more diverse and ecological set of regularities and inspect whether humans conform to our model's predictions also when the variety of regularities is much wider but could still fall into our proposed taxonomy.

Methods

Sequences' generative process. Each sequence was composed of 200 binary observations, that we refer to as A or B, and generated by one of the following 3 generative processes: (1) they could be random from the beginning to the end, (2) with a statistical bias introduced after an initial random part, (3) with a deterministic rule introduced after an initial random part.

Random parts. The random parts of sequences were generated by independent draws with equal probability for either binary outcome (i.e. as tosses of a fair coin). The resulting sequences contain on average as many As as Bs and as many repetitions (i.e. AA and BB) as alternations (i.e. AB and BA).

Change-points. In sequences entailing a regularity, a change-point separates the initial, random, part from the second, non-random, part. Unbeknownst to the subjects, the position of the change-point was drawn from a Gaussian distribution centred on the middle of the sequence (i.e. observation #100) with a standard deviation of 15 observations, truncated from observations #55 to #145 such that the change would not appear neither too early nor too late in the sequence. The resulting empirical distribution of change-points is centred on observation #100.65 ± 15.55 SD.

Statistical biases. Statistical regularities are obtained by drawing observations from a biased first-order Markov chain. Each statistical bias is thus fully described by the values of two first-order transition probabilities: $p(A|B) = 1 - p(B|B)$ and $p(B|A) = 1 - p(A|A)$. Note that unless otherwise specified, all transition probabilities are first-order. Importantly, transition probabilities also fully determine lower-order statistics: the frequency of items, $p(A) = 1 - p(B)$, and the frequency of alternation, $p(\text{alternation}) = 1 - p(\text{repetition})$ [18,52]. We used the following statistical biases (the numbers indicate $(p(A|B), p(B|A))$): more repetitions than alternations with (1/3, 1/3), (1/4, 1/4), and (1/5, 1/5); more alternations than repetitions with (2/3, 2/3), (3/4, 3/4) and (4/5, 4/5); more of

one item than the other with (1/3, 2/3), (1/4, 1/2) and (1/5, 4/5), which also bias the alternation frequency; or biased both in terms of alternation and item frequencies with (1/2, 1/4) and (3/4, 1/2). For repetition biases and alternation biases, the corresponding frequency of items is at chance level. We quantified the strength of the statistical biases as the Shannon entropy (i.e. H) of the distribution of pairs of items [77]:

$$H = - \sum_{X_1 \in \{A,B\}} \sum_{X_2 \in \{A,B\}} p\left(X_1 X_2\right) \cdot \log_2 p\left(X_1 X_2\right) \tag{1}$$

where the subscripts denote the position of the observation within the pair. The probability of a given pair of items can be computed using the generative transition probabilities:

$$p\left(X_1 X_2\right) = p\left(X_1\right) \cdot p\left(X_2 | X_1\right) \tag{2}$$

$$p\left(X\right) = \frac{p\left(X|Y\right)}{p\left(X|Y\right) + p\left(Y|X\right)} \tag{3}$$

Where X is either A or B and Y is either B or A. The entropy of our chosen statistical biases ranged from $H = 1.44$ to 1.92 bit, which contrasts with the entropy characterising random sequences, which is maximal, that is $H = 2$ bits. We used the limit $0 \cdot \log_2(0) = 0$ when it applied.

Deterministic rules. Deterministic rules were obtained by repeating a fixed pattern. Different pattern lengths were used: 4, 6, 8 and 10 observations. Moreover, patterns were chosen depending on the strength and type of bias in the apparent transition probabilities they induce. We term this bias "apparent", in opposition to "generative" since the deterministic rules are not generated according to a statistical bias, but indeed by repeating a pattern. Nonetheless, depending on the pattern used, a statistical bias may emerge: no bias with AABB and AAABBABB; a repetition bias with AAABBB, AAAABBBB and AAAAABBBBB; an alternation bias with AABABB, AABABABB and AABABABABB; and a bias in the frequency of items with AAABAB and AAABAABBAB. We selected several of these patterns, both in terms of bias type and bias strength so as to match with most of the biases characterising the statistical regularities we used. Apparent transition

probabilities can be computed for each deterministic rule (supposing an infinite number of repetitions of that pattern) and then used to estimate Shannon entropy, quantifying the strength of their apparent statistical bias using equation 1.

Experimental protocol. *Subjects.* A total of 28 subjects participated in the study. Data from 5 subjects were rejected because of very poor data quality reflecting a lack of comprehension of the instructions. Note, however, that our conclusions hold when using data from all the subjects; the reader can rerun our analyses with the code we provide online together with the full dataset. The data presented here come from 23 subjects (16 females) aged between 20 and 29 years old (mean age 23.91 ± 2.59 SD) from various education backgrounds including history of arts, law, translation, biology, engineering, psychology, etc. All subjects were right-handed. The study was approved by an independent ethics committee (CPP 08-021 Ile-de-France VII), and subjects gave their informed written consent before participating.

Stimulation. Stimuli consisted of two tones composed of three superimposed sine waves (350, 700, and 1400 Hz vs. 500, 1000, and 2000 Hz). The tones were 50 ms long, with 7 ms rise and fall times. They were randomly associated to A or B before each sequence. The stimulus-onset asynchrony was 300 ms long, thus resulting in 3.33 sounds per second. Stimulation was delivered using MATLAB and Psychtoolbox [78].

Triangular arena. The triangular arena is an equilateral triangle displayed on a touch screen and whose vertices correspond to the 3 possible generative hypotheses: the bottom vertex corresponds to a random process and the two upper vertices correspond to non-random processes (i.e. statistics and rules) whose respective side (i.e. left or right) was counterbalanced across subjects. For the sake of averaging and clarity of reporting, the trajectories of half of the subjects were thus vertically mirrored such that for all subjects the statistical bias hypothesis corresponds to the left upper vertex

and the deterministic rule hypothesis to the right upper vertex. A point in that triangular space can be converted into posterior probabilities for each generative hypothesis (by means of barycentric coordinates summing to 1, see below). The centre of the triangle corresponds to equal probabilities (i.e. (1/3, 1/3, 1/3)). Because all sequences start with a random part, observers started at the bottom vertex and updated their locations in the triangle as observations were delivered in a way that reflects their estimates of the probabilities for the three possible hypotheses given the received observations. Subjects were allowed to move freely inside the triangular arena: they could pause at a particular location of the triangle or change their mind for instance by returning near the random vertex after a suspected regularity was discarded. The triangular arena (9.7 cm wide × 8.4 cm high, 800 pixels wide × 693 pixels high) was displayed using MATLAB and Psychtoolbox [78] on a 14 inches touch screen (31 cm wide × 17.4 cm high, 2560 pixels wide × 1440 pixels high) computer (HP Pavilion x360) that was lying horizontally on a table. Finger position was collected using MATLAB immediately after each auditory stimulus was played, thereby resulting in a trajectory composed of 200 pairs of cartesian coordinates (a_k and b_k). When the recorded finger position was outside of the triangle (because of subject's motor imprecision), the positions were projected back to the nearest point on the triangle border (through Euclidean distance minimization, average distance = 5.83 pixels ± 3.88 SD, [1.77, 16.01]).

Training. Subjects were first given written instructions. They were then presented with several short example sequences (70 observations, no change-point) and told the underlying generative process in order for them to understand the differences between the two types of regularities and the random process. The experimenter provided examples of how to use the triangular arena to report probabilistic beliefs with 4 example sequences (1 random, 2 random-to-statistical and 1 random-to-rule). Finally, subjects practiced on 7 example sequences (2 random, 3 random-to-statistical and 2 random-to-rule). The experimenter then answered possible questions and launched the experiment. The regularities used during training were different from those used in the experiment.

Experiment. During the experiment, subjects were presented with 33 sequences (10 random, 13 random-to-statistical and 10 random-to-rule) whose order was randomly defined for each subject. These sequences were allocated to 4 different sessions (each composed of 8 or 9 sequences) interspersed with pauses. Two of the random-to-statistical sequences were duplicates (of (1/2, 1/4) and (3/4, 1/2)) that were added in order to assess within-subject variability in stochastic conditions. They are simply discarded from the analyses presented here. All trials follow the same structure. Subjects initiated the trial by touching a button on the touch screen. The triangular arena was then displayed on the screen and a 3-second countdown was initiated. During the following minute, the sequence was played, and subjects moved their finger inside the triangular arena. Once the presentation of the sequence finished, subjects were asked several post-sequence questions. They were first asked whether they thought retrospectively that the sequence entailed a regularity or not (detection). In the absence of noticed regularities, the trial was terminated. Otherwise, they were asked several other questions. Firstly, which type of regularity (statistics vs. rule) occurred (discrimination)? Secondly, when the change-point occurred? In that case, they used a continuous horizontal scale from observation #1 to #200 and also rated their confidence in that estimate using a continuous vertical scale. Those two questions are in retrospect, which is in contrast with the third question in which they were asked to estimate when they had realized, during the sequence presentation, that a change-point had occurred. In that case, they also used a continuous horizontal scale from observation reported as change-point (at the second question) to observation #200 and also rated their confidence in that estimate using a continuous vertical scale. The data from the third question were not used here.

Normative two-system model. We derive the ideal observer model for the task, which is the Bayes-optimal solution using the actual task structure [79,80]. Comparing subjects behaviour against this benchmark inference therefore allows the identification of signatures of a normative inference in

human behaviour [49,50]. Practically speaking, the model is presented with a given sequence y, of length $|y| = K$, and returns the posterior probability of each possible hypothesis $\mathcal{H}_h$ (i.e. random, random-to-statistical and random-to-rule processes) using Bayes' theorem, as the sequence unfolds in time:

$$p\left(\mathcal{H}_h|y\right) = \frac{p\left(y|\mathcal{H}_h\right) \cdot p\left(\mathcal{H}_h\right)}{\sum_h p\left(y|\mathcal{H}_h\right) \cdot p\left(\mathcal{H}_h\right)} \tag{4}$$

The assumptions of the model closely correspond to the task structure, and they determine its prior and likelihood functions. The likelihood of the sequence under each hypothesis $p\left(y|\mathcal{H}_h\right)$ is derived below. The prior probability over hypotheses was uniform, such that $p\left(\mathcal{H}_h\right) = \frac{1}{3}$. Non-uniform prior distributions over hypotheses have been explored but fail to significantly better explain subjects' behaviour compared to the simpler uniform prior distribution.

Sequence likelihood under a random hypothesis. Under the random hypothesis $\mathcal{H}_{\mathrm{rand}}$, all observations are at chance level and independent. The likelihood of a sequence is thus the product of the chance-level likelihoods of each observation. Because the observations are binary, the process amounts to independent tosses of a fair coin, and the likelihood of each observation is ½:

$$p\left(y|\mathcal{H}_{\mathrm{rand}}\right) = p\left(\mathrm{A}|\mathcal{H}_{\mathrm{rand}}\right)^{N_y^{\mathrm{A}}} \cdot p\left(\mathrm{B}|\mathcal{H}_{\mathrm{rand}}\right)^{N_y^{\mathrm{B}}} = \left(\frac{1}{2}\right)^K \tag{5}$$

Where N_y^{X} denotes the number of times X has been observed in the sequence y.

Likelihood of a sequence under non-random hypotheses. For the remaining two hypotheses, a change-point delineates the initial, random, part of the sequence from the second, non-random, part. However, the location of the change-point is unknown and assumed random. In order to get rid of this unknown factor, one must therefore consider all possible j_k positions of the change-point (i.e. after the 1st, 2nd, …, up to the $(K-1)$th observation), and marginalize over these positions:

$$p\left(y|\mathcal{H}_{\mathrm{rand,reg}}\right) = \sum_{k=1}^{K-1} p\left(y|j_k, \mathcal{H}_{\mathrm{rand,reg}}\right) \cdot p\left(j_k\right) \tag{6}$$

The likelihood of the sequence is, for a given change-point location, the product of the likelihood of the random part and the non-random part of the sequence:

$$p\left(y|j_k, \mathcal{H}_{\mathrm{rand,reg}}\right) = p\left(y_{1:k}|\mathcal{H}_{\mathrm{rand}}\right) \cdot p\left(y_{k+1:K}|\mathcal{H}_{\mathrm{reg}}\right) \tag{7}$$

The likelihood of the first part of the sequence ($y_{1:k}$, noted below y_{rand} for brevity) is the likelihood of a random sequence (equation 5), the likelihood of the second part ($y_{k+1:K}$, noted below y_{stat} or y_{rule} for brevity, whose length is $|y_{\mathrm{stat\ or\ rule}}| = K - j_k$) depends upon the type of regularity that is considered (either statistical bias or deterministic rule, see below). The prior distribution over change-point positions is set to be uniform, such that all positions are a priori equally likely $p\left(j_k\right) = \frac{1}{200}$.

Sequence likelihood under a statistical bias hypothesis. The random-to-statistical hypothesis $p\left(y|\mathcal{H}_{\mathrm{rand,stat}}\right)$ considers that, after the change-point, the sequence likelihood depends on a matrix of transition probabilities between successive observations characterising a first-order Markov chain [52]. The generative matrix of transition probabilities is unknown and random from one sequence to the next. In order to get rid of this unknown factor, one must therefore consider all possible values of transition probabilities $\boldsymbol{\theta}$ and marginalize over them. The likelihood of the second, statistical, part of the sequence y_{stat} is thus:

$$p\left(y_{\mathrm{stat}}|\mathcal{H}_{\mathrm{stat}}\right) = \int_0^1 p\left(y_{\mathrm{stat}}|\boldsymbol{\theta}, \mathcal{H}_{\mathrm{stat}}\right) \cdot p\left(\boldsymbol{\theta}|\mathcal{H}_{\mathrm{stat}}\right) \mathrm{d}\boldsymbol{\theta} \tag{8}$$

For a model that estimates transition probabilities between consecutive stimuli, the likelihood of a given observation depends only on the estimated transition probabilities and the previous stimulus. For simplicity, the first observation can be considered as arbitrary such that its likelihood equals chance level. This results in:

$$p\left(y_{\mathrm{stat}}|\boldsymbol{\theta}, \mathcal{H}_{\mathrm{stat}}\right) = p\left(y_1^{\mathrm{stat}}|\boldsymbol{\theta}\right) \cdot \prod_{k=2}^{|y_{\mathrm{stat}}|} p\left(y_k^{\mathrm{stat}}|y_{k-1}^{\mathrm{stat}}, \boldsymbol{\theta}, \mathcal{H}_{\mathrm{stat}}\right)$$

$$= \frac{1}{2} \cdot \left(\theta_{\mathrm{A|B}}^{N_{y_{\mathrm{stat}}}^{\mathrm{A|B}}} \cdot \left(1 - \theta_{\mathrm{A|B}}\right)^{N_{y_{\mathrm{stat}}}^{\mathrm{B|B}}}\right) \cdot \left(\theta_{\mathrm{B|A}}^{N_{y_{\mathrm{stat}}}^{\mathrm{B|A}}} \cdot \left(1 - \theta_{\mathrm{B|A}}\right)^{N_{y_{\mathrm{stat}}}^{\mathrm{A|A}}}\right) \quad (9)$$

Where $\boldsymbol{\theta}$ is a vector of two transition probabilities $\boldsymbol{\theta} = [\theta_{\mathrm{A|B}}, \theta_{\mathrm{B|A}}]$ and $N^{\mathrm{X|Y}}$ denotes the number of YX pairs in the statistical part of the sequence. Equation 9 corresponds to the product of two Beta distributions, which parameters are the transition counts plus one:

$$p\left(y_{\mathrm{stat}}|\mathcal{H}_{\mathrm{stat}}\right) = \frac{1}{2} \cdot \int_0^1 \mathrm{Beta}\left(\theta_{\mathrm{A|B}}|N_{\mathrm{A|B}} + 1, N_{\mathrm{B|B}} + 1\right) \cdot \mathrm{Beta}\left(\theta_{\mathrm{B|A}}|N_{\mathrm{B|A}} + 1, N_{\mathrm{A|A}} + 1\right) \mathrm{d}\boldsymbol{\theta} \quad (10)$$

The integral over such Beta distributions has analytical solutions that can be used to compute sequence likelihood in equation 8 when a uniform conjugate prior over $\boldsymbol{\theta}$ is used [81].

Sequence likelihood under a deterministic rule hypothesis. The random-to-rule hypothesis $p\left(y|\mathcal{H}_{\mathrm{rand,rule}}\right)$ considers that, after the change-point, the sequence can be described as the repetition of a particular pattern. However, this pattern is unknown. In order to get rid of this unknown factor, one must therefore consider all possible patterns, denoted $\{\mathcal{R}\}$, and marginalize over them. The likelihood of the second, deterministic, part of the sequences y_{rule} is thus:

$$p\left(y_{\mathrm{rule}}|\mathcal{H}_{\mathrm{rule}}\right) = \sum_{r \in \{\mathcal{R}\}} p\left(y_{\mathrm{rule}}|r, \mathcal{H}_{\mathrm{rule}}\right) \cdot p\left(r|\mathcal{H}_{\mathrm{rule}}\right) \quad (11)$$

We consider that the set of all patterns $\{\mathcal{R}\}$ can be split in two disjoint subsets. The first one corresponds to all patterns whose length is equal or smaller than the length of the deterministic sequence: this is the set of patterns potentially entirely observed in the current sequence $\{\mathcal{R}^{\mathrm{o}}\}$. The second subset corresponds to patterns whose length is larger than the length of the deterministic sequence: this is the set of partially observed patterns $\{\mathcal{R}^{\mathrm{u}}\}$. The size of that latter set depends upon the maximum pattern length allowed, denoted l. We used $l = 10$ observations such that the model can detect all the patterns used in the experiment while remaining psychologically plausible given human memory limits. Conveniently, this parameter choice does not significantly impact the

predictions of the model for $l > 10$ as it allows the detection of all patterns used here (*Supplementary Fig. 3*). For both patterns' sets, the likelihood of a particular pattern r given the observed sequence is all or none:

$$p\left(y_{\text{rule}}|r\right) = \begin{cases} 1 & \text{if } y_{\text{rule}} \text{ is equivalent to the repetition of } r \\ 0 & \text{otherwise} \end{cases} \tag{12}$$

Among $\{\mathcal{R}^{\text{o}}\}$, there is at most one (and thus possibility zero) pattern of each given length that is compatible with the observed sequence (i.e. for which the sequence likelihood is not null). Possible extensions of this pattern found in $\{\mathcal{R}^{\text{u}}\}$ also result in a sequence likelihood equal to one, and all the others patterns are incompatible with the observed sequence (and thus result in a likelihood that is null).

The prior probability of patterns reflects an incremental procedure used to generate the patterns: at each iteration, the current pattern is either selected, which terminates the procedure, or extended with observation A or B, with equal probability among those 3 possibilities. The process terminates when the maximum pattern length l is reached. With such a process, longer patterns are less probable than shorter ones, the prior therefore favours shorter patterns (see first ratio in equation 13), following the "size-principle" which has been widely used in the past [82]. For a given pattern r of length $|r|$, the prior probability is:

$$p\left(r\right) = \frac{1}{3^{|r|-1}} \cdot \frac{1}{\frac{9}{2} \cdot \left(1 - \left(\frac{2}{3}\right)^{l+1} - \frac{3}{2}\right)} \tag{13}$$

Where the second ratio ensures, given that l is the maximum pattern length allowed, that $\sum_{r \in \{\mathcal{R}\}} p\left(r\right) = 1$.

Posterior distribution of change-point position. Besides computing the posterior probability of the three different generative processes given an input sequence, it is also possible to compute the posterior probability of the change-point being at one particular position, for the two hypotheses that assume the existence of such a change-point:

$$\forall h \in \{\text{stat}, \text{rule}\}, \forall j_k \in \{1, 2, \ldots, 200\} : p\left(j_k | y, \mathcal{H}_{\text{rand},h}\right) = \frac{p\left(y | j_k, \mathcal{H}_{\text{rand},h}\right) \cdot p\left(j_k\right)}{\sum_{k=1}^{200} p\left(y | j_k, \mathcal{H}_{\text{rand},h}\right) \cdot p\left(j_k\right)} \tag{14}$$

The likelihood of an observed sequence containing a change-point has been previously defined in equation 7. By contrast, if the change-point is in the yet unobserved part of the sequence, then the current sequence is still necessarily random:

$$p\left(y | j_k, \mathcal{H}_{\text{rand},h}\right) = \begin{cases} p\left(y_{\text{rand}} | \mathcal{H}_{\text{rand}}\right) \cdot p\left(y_h | \mathcal{H}_h\right) & \text{if } j_k \in \{1, 2, \ldots, K\} \\ p\left(y | \mathcal{H}_{\text{rand}}\right) = \left(\frac{1}{2}\right)^K & \text{if } j_k \in \{K+1, K+2, \ldots, 200\} \end{cases} \tag{15}$$

Those hypothesis-specific posterior distributions over change-points can then be combined by means of Bayesian model averaging:

$$\forall j_k \in \{1, 2, \ldots, 200\} : p\left(j_k | y\right) = \frac{\sum_{h \in \{\text{stat},\text{rule}\}} p\left(j_k | y, \mathcal{H}_{\text{rand},h}\right) \cdot p\left(\mathcal{H}_{\text{rand},h} | y\right)}{\sum_{h \in \{\text{stat},\text{rule}\}} p\left(\mathcal{H}_{\text{rand},h} | y\right)} \tag{16}$$

The most likely position of the change-point is therefore the maximum a posteriori of this distribution and the confidence related to that estimate can be assessed by measuring the log-precision of the posterior distribution [19,44,45]:

$$\log \text{precision} = \log\left(\frac{1}{\text{variance}}\right) = -\log\left(\sum_{k=1}^{200} p\left(j_k | y\right) \cdot k^2 - \sum_{k=1}^{200} p\left(j_k | y\right) \cdot k\right) \tag{17}$$

Normative single-system model. An alternative to the dichotomy between the statistical bias hypothesis (with order 1 transition probabilities) and the deterministic rule hypothesis (with patterns) is that subjects monitor order 1 vs. higher-order transition probabilities. To explore this possibility, we replaced the deterministic hypothesis $\mathcal{H}_{\text{rule}}$ (with patterns) in the *normative two-system model* (all other computations remain the same) with a higher-order statistical bias hypothesis $\mathcal{H}'_{\text{rule}}$, with Markov chains of various orders, also referred to as *n-grams* models in the linguistic literature [83]. In those cases, the likelihood of the second part of the sequence becomes:

$$p\left(y_{\text{rule}} | \mathcal{H}'_{\text{rule}}\right) = \frac{1}{2^\alpha} \cdot \int_0^1 \prod_{i=1}^{2^\alpha} \text{Beta}\left(\theta_{A|T_i} | N_{A|T_i} + 1 - d, N_{B|T_i} + 1 - d\right) \mathrm{d}\boldsymbol{\theta} \tag{18}$$

Where α is the order of the chain, resulting in 2^{α} transitions $\mathbf{T}$. For instance, a Markov chain of order $\alpha = 1$ corresponds to the transitions $\mathbf{T} = [X|A, X|B]$, a chain of order $\alpha = 2$ considers the transitions $\mathbf{T} = [X|AA, X|AB, X|BA, X|BB]$, etc. In addition to varying the order of the Markov chain, we also varied a parameter d that controls the preference for predictability of the conjugate prior. For instance, $d = 0$ corresponds to a uniform prior: the next observation can have any likelihood, from 0 to 1 with equal probability; while $d > 0$ favours extreme probabilities: in this case, the observer expects more extreme probabilities (toward $p = 0$ and 1) to be more likely than intermediate probabilities (around $p = \frac{1}{2}$) [84]. The likelihood function we used for the statistical bias hypothesis (equation 10) is thus a spatial case of equation 18 where $\alpha = 1$ and $d = 0$, that is a first-order Markov chain with a uniform prior.

Non-commensurate two-system model. An alternative to the rational arbitration between the non-random hypotheses by means of a common probabilistic currency is that subjects evaluate the posterior probability of each non-random hypothesis (against the random hypothesis) independently of one another. To explore this possibility, and contrary to the *normative two-system model*, we computed posterior probability of each non-random hypothesis separately, given the observed sequence:

$$\forall h \in \{\text{stat}, \text{rule}\} : q\left(\mathcal{H}_h | y\right) = \frac{p\left(y|\mathcal{H}_h\right) \cdot p\left(\mathcal{H}_h\right)}{p\left(y|\mathcal{H}_h\right) \cdot p\left(\mathcal{H}_h\right) + p\left(y|\mathcal{H}_{\text{rand}}\right) \cdot p\left(\mathcal{H}_{\text{rand}}\right)} \tag{19}$$

Where the prior probabilities are set to chance level (i.e. $\frac{1}{2}$ in that case). From equation 19, it follows that the probability of the random hypothesis depends upon the non-random hypothesis that is considered:

$$\forall h \in \{\text{stat}, \text{rule}\} : q\left(\mathcal{H}_{\text{rand}} | y, \mathcal{H}_h\right) = 1 - q\left(\mathcal{H}_h | y\right) \tag{20}$$

In order to report the probabilities of the 3 hypotheses in the triangular arena (which requires they sum to 1), we used four different combination functions: (1) linear difference, (2) max rule (taking the best hypothesis and discarding the other), (3) a sigmoid based on the difference between the

posterior probability of non-random hypotheses, and (4) a sigmoid based on the log-ratio between the posterior probability of non-random hypotheses. For the two latter cases, one parameter (β) controls the slope of the sigmoid function (we explored 100 values equally spaced logarithmically between $10^{0.5}$ and $10^{9.5}$).

$$\text{Linear difference:} \quad w_{\text{rule}} = \frac{q\left(\mathcal{H}_{\text{rule}}|y\right)}{q\left(\mathcal{H}_{\text{rule}}|y\right) + q\left(\mathcal{H}_{\text{stat}}|y\right)} \tag{21}$$

$$\text{Max rule:} \quad w_{\text{rule}} = \begin{cases} 0 \text{ if } q\left(\mathcal{H}_{\text{rule}}|y\right) < q\left(\mathcal{H}_{\text{stat}}|y\right) \\ \frac{1}{2} \text{ if } q\left(\mathcal{H}_{\text{rule}}|y\right) = q\left(\mathcal{H}_{\text{stat}}|y\right) \\ 1 \text{ if } q\left(\mathcal{H}_{\text{rule}}|y\right) > q\left(\mathcal{H}_{\text{stat}}|y\right) \end{cases} \tag{22}$$

$$\text{Sigmoid on the difference:} \quad w_{\text{rule}} = \frac{1}{1 + \exp\left(-\beta \cdot \left(q\left(\mathcal{H}_{\text{rule}}|y\right) - q\left(\mathcal{H}_{\text{stat}}|y\right)\right)\right)} \tag{23}$$

$$\text{Sigmoid on the log-ratio:} \quad w_{\text{rule}} = \frac{1}{1 + \exp\left(-\beta \cdot \log \frac{q(\mathcal{H}_{\text{rule}}|y)}{q(\mathcal{H}_{\text{stat}}|y)}\right)} \tag{24}$$

The pseudo posterior probability of the random hypothesis is first obtained by a weighted combination of the probability of the random hypothesis under each non-random hypothesis:

$$p\left(\mathcal{H}_{\text{rand}}|y\right) = \left(1 - w_{\text{rule}}\right) \cdot q\left(\mathcal{H}_{\text{rand}}|y, \mathcal{H}_{\text{stat}}\right) + w_{\text{rule}} \cdot q\left(\mathcal{H}_{\text{rand}}|y, \mathcal{H}_{\text{rule}}\right) \tag{25}$$

Then, the pseudo posterior probabilities of the non-random hypotheses are computed based on the weights from the previously mentioned combination functions which guarantees that the probabilities over the 3 hypotheses sum to 1:

$$\begin{aligned} p\left(\mathcal{H}_{\text{rule}}|y\right) &= \left(1 - p\left(\mathcal{H}_{\text{rand}}|y\right)\right) \cdot w_{\text{rule}} \\ p\left(\mathcal{H}_{\text{stat}}|y\right) &= \left(1 - p\left(\mathcal{H}_{\text{rand}}|y\right)\right) \cdot \left(1 - w_{\text{rule}}\right) \end{aligned} \tag{26}$$

From finger position to posterior probability. A given pair of cartesian coordinates (a, b) specifying the subject's finger position can be turned into the posterior probability of each hypothesis $p\left(\mathcal{H}_h|y\right)$ by performing a conversion to the barycentric coordinate system using the following equations:

$$\begin{cases} p\left(\mathcal{H}_{\mathrm{stat}}|y\right) = \dfrac{(b_{\mathrm{stat}} - b_{\mathrm{rand}}) \cdot (a - a_{\mathrm{rand}}) + (a_{\mathrm{rand}} - a_{\mathrm{rule}}) \cdot (b - b_{\mathrm{rand}})}{(b_{\mathrm{stat}} - b_{\mathrm{rand}}) \cdot (a_{\mathrm{stat}} - a_{\mathrm{rand}}) + (a_{\mathrm{rand}} - a_{\mathrm{rule}}) \cdot (b_{\mathrm{stat}} - b_{\mathrm{rand}})} \\ p\left(\mathcal{H}_{\mathrm{rule}}|y\right) = \dfrac{(b_{\mathrm{rand}} - b_{\mathrm{stat}}) \cdot (a - a_{\mathrm{rand}}) + (a_{\mathrm{stat}} - a_{\mathrm{rand}}) \cdot (b - b_{\mathrm{rand}})}{(b_{\mathrm{rule}} - b_{\mathrm{rand}}) \cdot (a_{\mathrm{stat}} - a_{\mathrm{rand}}) + (a_{\mathrm{rand}} - a_{\mathrm{stat}}) \cdot (b_{\mathrm{stat}} - b_{\mathrm{rand}})} \\ p\left(\mathcal{H}_{\mathrm{rand}}|y\right) = 1 - p\left(\mathcal{H}_{\mathrm{stat}}|y\right) - p\left(\mathcal{H}_{\mathrm{rule}}|y\right) \end{cases} \tag{27}$$

Where a_h is the horizontal position of hypothesis h on the screen and b_h its vertical position (both in pixels). Conversely, the posterior probability of the three hypotheses estimated by the model can be reported on the touch screen by applying a barycentric-to-cartesian transformation:

$$\begin{cases} a = a_{\mathrm{stat}} \cdot p\left(\mathcal{H}_{\mathrm{stat}}|y\right) + a_{\mathrm{rule}} \cdot p\left(\mathcal{H}_{\mathrm{rule}}|y\right) + a_{\mathrm{rand}} \cdot p\left(\mathcal{H}_{\mathrm{rand}}|y\right) \\ b = b_{\mathrm{stat}} \cdot p\left(\mathcal{H}_{\mathrm{stat}}|y\right) + b_{\mathrm{rule}} \cdot p\left(\mathcal{H}_{\mathrm{rule}}|y\right) + b_{\mathrm{rand}} \cdot p\left(\mathcal{H}_{\mathrm{rand}}|y\right) \end{cases} \tag{28}$$

Triangular histograms of positions in the triangle. The histograms were obtained by segmenting the triangle into 87 vertical and 101 horizontal rectangular parcels (since the height of an equilateral triangle is 0.87 the size of its edge) and counting the frequency of occupancy of each of these parcels. Before counting, the subjects' and model's trajectories were linearly interpolated (resulting in 3 times more data points). The resulting 2D histogram was then smoothed with a Gaussian kernel ($\sigma = 4$) and normalized such that the sum over parcels equals 1.

Measures of dynamics. We used different measures to quantify different aspects of the inference dynamics.

Detection-point. Detection-points are defined as the number of observations after the change-point for the posterior probability in the correct hypothesis to become likely than any of the others (i.e. larger than ½).

Detection slope. A sigmoid function was used to approximate the slope of detection dynamics:

$$p^* = t_u \cdot \frac{1}{1 + \exp\left(-s \cdot k - k_0\right)} + t_b \tag{29}$$

Where s, k_0, t_u and t_b are the 4 free parameters of the function (respectively the slope, the intercept and the upper and lower bounds) that transforms the position k of an observation in the sequence into a given probability estimate p^*. They were fitted using a grid search technique that finds the set of 4 parameters that minimizes the MSE along the following parameter grid: for the slope parameter, 31 values logarithmically equally spaced between decades 10^{-3} and 10^{0}; for the intercept parameter, any observation number between 0 (change-point) and 200; for the offset parameter, 11 values linearly spaced between 0 and 0.5 with a step of 0.05; for the scaling parameter, 11 value linearly spaced between 0.5 and 1 with a step of 0.05.

Belief update. The amount of belief update quantifies the average increase in posterior probability each observation in the non-random part of the sequence provided to the correct non-random hypothesis.

$$\text{belief update} = \frac{p\left(\mathcal{H}_{\text{correct}}|y_{1:200}\right) - p\left(\mathcal{H}_{\text{correct}}|y_{1:\text{change}}\right)}{200 - \text{change}} = \left\langle \frac{\mathrm{d}p\left(\mathcal{H}_{\text{correct}}|y_{\text{change}:200}\right)}{\mathrm{d}y} \right\rangle_k \tag{30}$$

Belief difference. An index reflecting the competition among the non-random hypotheses (statistical bias vs. deterministic rule) independently from the posterior probability of the random hypothesis was also computed:

$$\text{belief difference} = \frac{p\left(\mathcal{H}_{\text{rule}}|y\right) - p\left(\mathcal{H}_{\text{stat}}|y\right)}{p\left(\mathcal{H}_{\text{rule}}|y\right) + p\left(\mathcal{H}_{\text{stat}}|y\right)} \tag{31}$$

This index ranges from –1 (statistical bias) to 1 (deterministic rule).

Detection abruptness. The abruptness index quantifies how abrupt is the update of a given hypothesis by averaging the absolute value of its second order derivative over observations after the change-point:

$$\text{abruptness} = \left\langle \left| \frac{\mathrm{d}^2 p \left(\mathcal{H}_{\text{rule}} | y \right)}{\mathrm{d}y^2} \right| \right\rangle_{\text{change:200}} \tag{32}$$

It is therefore higher for frequently interrupted dynamics compared to continuously updated ones. We computed this index for each pattern for both subjects and models and use it to compute the mean squared error (MSE, averaged over patterns) of each model.

Signal detection theory. The detection sensitivity index $d'_{\text{detect.}}$ between random and non-random sequences is computed by comparing non-random sequences accurately labelled as non-random vs. random sequences erroneously reported as non-random using answers to the first post-sequence question (presence of regularity?):

$$
\begin{aligned}
d'^{\text{stat}}_{\text{detec}} &= \Phi \left[p \left(\text{reg} | \text{rand-to-stat} \right) \right] - \Phi \left[p \left(\text{reg} | \text{rand} \right) \right] \\
d'^{\text{rule}}_{\text{detec}} &= \Phi \left[p \left(\text{reg} | \text{rand-to-rule} \right) \right] - \Phi \left[p \left(\text{reg} | \text{rand} \right) \right]
\end{aligned}
\tag{33}
$$

Where Φ is the inverse cumulative function of the Gaussian distribution (with $\mu = 0$ and $\sigma = 1$). By contrast, the discrimination sensitivity index $d'_{\text{discr.}}$ between the two random-to-non-random sequences is computed by comparing correct to incorrect categorisations of regularities at the second post-sequence question (type of regularity?):

$$
\begin{aligned}
d'_{\text{discr}} &= \frac{1}{\sqrt{2}} \cdot \left(\Phi \left[p \left(\text{stat} | \text{stat} \right) \right] - \Phi \left[p \left(\text{stat} | \text{rule} \right) \right] \right) \\
&= \frac{1}{\sqrt{2}} \cdot \left(\Phi \left[p \left(\text{rule} | \text{rule} \right) \right] - \Phi \left[p \left(\text{rule} | \text{stat} \right) \right] \right)
\end{aligned}
\tag{34}
$$

The $\frac{1}{\sqrt{2}}$ factor guarantees that if both regularities were to lie on orthogonal dimensions at a given distance from the random hypothesis (the origin), then, according to Pythagora's theorem, the distance between the non-random distributions should be separated by a distance larger than $\sqrt{2}$ the initial distance [85]. In order to be able to compare detection and discrimination d', the latter is therefore divided by that same amount. Note that we do not hypothesize such orthogonal arrangement; instead, the orthogonality between the two dimensions is estimated (using Al-Kashi's theorem) from the 3 sensitivity indices that define the length of the triangles' edges. The same

sensitivity measures are computed for the model which identifies the generative process based on the maximum a posteriori probability at the end of the sequence.

Statistics. All dispersion metrics reported in the figures are standard errors of the mean (s.e.m., with $n = 23$ subjects), except shaded areas around regression lines which reflect 95% confidence intervals around estimated regression coefficients. Reported correlation coefficients (noted r) are Pearson coefficients. Parametric frequentist statistics with repeated measures are performed. For the t-tests, the mean value of the paired difference, 95% confidence intervals, an estimate of effect size (Cohen's d), the t statistic, the corresponding degrees of freedom (noted t_{df}) and p-values (noted p) are reported. For the ANOVAs, an estimate of effect size (w^2), the value of the F statistic, the corresponding degrees of freedom (noted F_{df}) and p-values (noted p) are reported. All statistical tests were two-tailed at with a type I error risk of 0.05. Bayesian t-tests were used to quantify the evidence in favour of the null hypothesis using a scale factor of 0.707 [86].

Code availability. The MATLAB code used to run simulations of the different models, analyse the results and reproduce all the figures is available on *GitHub* (https://github.com/maheump/Emergence).

Data availability. The dataset presented in the current study is available on *GitHub* (https://github.com/maheump/Emergence).

References

1. Clegg, B. A., Digirolamo, G. J. & Keele, S. W. Sequence learning. *Trends Cogn. Sci.* **2**, 275–281 (1998).

2. Dehaene, S., Meyniel, F., Wacongne, C., Wang, L. & Pallier, C. The neural representation of

sequences: From transition probabilities to algebraic patterns and linguistic trees. *Neuron* **88**, 2–19 (2015).

3. Lashley, K. S. *The Problem of Serial Order in Behavior*. (1951).

4. Friston, K. A theory of cortical responses. *Philos. Trans. R. Soc. B Biol. Sci.* **360**, 815–836 (2005).

5. Rao, R. P. & Ballard, D. H. Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* **2**, 79–87 (1999).

6. Spratling, M. W. A review of predictive coding algorithms. *Brain Cogn.* **112**, 92–97 (2017).

7. Friston, K. The free-energy principle: A unified brain theory? *Nat. Rev. Neurosci.* **11**, 127–138 (2010).

8. Friston, K., Rosch, R., Parr, T., Price, C. & Bowman, H. Deep temporal models and active inference. *Neurosci. Biobehav. Rev.* **90**, 486–501 (2018).

9. Summerfield, C. & de Lange, F. P. Expectation in perceptual decision making: Neural and computational mechanisms. *Nat. Rev. Neurosci.* **15**, 745–756 (2014).

10. Barascud, N., Pearce, M. T., Griffiths, T. D., Friston, K. J. & Chait, M. Brain responses in humans reveal ideal observer-like sensitivity to complex acoustic patterns. *Proc. Natl. Acad. Sci. U. S. A.* **113**, E616–25 (2016).

11. Barczak, A. *et al.* Top-down, contextual entrainment of neuronal oscillations in the auditory thalamocortical circuit. *Proc. Natl. Acad. Sci.* **115**, E7605–E7614 (2018).

12. Boubenec, Y., Lawlor, J., Górska, U., Shamma, S. & Englitz, B. Detecting changes in dynamic and complex acoustic environments. *eLife* **6**, 1929 (2017).

13. Herrmann, B. & Johnsrude, I. S. Neural signatures of the processing of temporal patterns in sound. *J. Neurosci.* **38**, 5466–5477 (2018).

14. Skerritt-Davis, B. & Elhilali, M. Detecting change in stochastic sound sequences. *PLoS Comput. Biol.* **14**, e1006162 (2018).

15. Brown, S. D. & Steyvers, M. Detecting and predicting changes. *Cognit. Psychol.* **58**, 49–67

(2009).

16. Gallistel, C. R., Krishan, M., Liu, Y., Miller, R. & Latham, P. E. The perception of probability. *Psychol. Rev.* **121**, 96–123 (2014).

17. Glaze, C. M., Kable, J. W. & Gold, J. I. Normative evidence accumulation in unpredictable environments. *eLife* **4**, (2015).

18. Maheu, M., Dehaene, S. & Meyniel, F. Brain signatures of a multiscale process of sequence learning in humans. *eLife* **8**, 275 (2019).

19. Meyniel, F., Schlunegger, D. & Dehaene, S. The sense of confidence during probabilistic learning: A normative account. *PLoS Comput. Biol.* **11**, e1004305 (2015).

20. Nassar, M. R., Wilson, R. C., Heasly, B. & Gold, J. I. An approximately Bayesian delta-rule model explains the dynamics of belief updating in a changing environment. *J. Neurosci.* **30**, 12366–12378 (2010).

21. Saffran, J. R., Aslin, R. N. & Newport, E. L. Statistical learning by 8-month-old infants. *Science* **274**, 1926–1928 (1996).

22. Glascher, J., Daw, N. D., Dayan, P. & O'Doherty, J. P. States versus rewards: Dissociable neural prediction error signals underlying model-based and model-free reinforcement learning. *Neuron* **66**, 585–595 (2010).

23. Gold, J. I. & Shadlen, M. N. The neural basis of decision making. *Annu. Rev. Neurosci.* **30**, 535–574 (2007).

24. Hollerman, J. R. & Schultz, W. Dopamine neurons report an error in the temporal prediction of reward during learning. *Nat. Neurosci.* **1**, 304–309 (1998).

25. Huys, Q. J. M. *et al.* Bonsai trees in your head: How the pavlovian system sculpts goal-directed choices by pruning decision trees. *PLoS Comput. Biol.* **8**, e1002410 (2012).

26. Lee, D., Seo, H. & Jung, M. W. Neural basis of reinforcement learning and decision making. *Annu. Rev. Neurosci.* **35**, 287–308 (2012).

27. Rescorla, R. & Wagner, A. *A Theory of Pavlovian Conditioning: Variations in the*

Effectiveness of Reinforcement and Nonreinforcement. (Appletone-Century-Crofts, New York, 1972).

28. Sutton, R. S. & Barto, A. G. *Reinforcement learning*. (MIT Press, 1998). doi:10.5772/5275.

29. Amalric, M. *et al.* The language of geometry: Fast comprehension of geometrical primitives and rules in human adults and preschoolers. *PLoS Comput. Biol.* **13**, e1005273 (2017).

30. Friederici, A. D., Bahlmann, J. r, Heim, S., Schubotz, R. I. & Anwander, A. The brain differentiates human and non-human grammars: Functional localization and structural connectivity. *Proc. Natl. Acad. Sci.* **103**, 2458–2463 (2006).

31. Konovalov, A. & Krajbich, I. Neurocomputational dynamics of sequence learning. *Neuron* **98**, 1282–1293.e4 (2018).

32. Landmann, C. *et al.* Dynamics of prefrontal and cingulate activity during a reward-based logical deduction task. *Cereb. Cortex* **17**, 749–759 (2007).

33. Lieberman, M. D., Chang, G. Y., Chiao, J., Bookheimer, S. Y. & Knowlton, B. J. An event-related fMRI study of artificial grammar learning in a balanced chunk strength design. *J. Cogn. Neurosci.* **16**, 427–438 (2004).

34. Marcus, G. F., Vijayan, S., Rao, S. B. & Vishton, P. M. Rule learning by seven-month-old infants. *Science* **283**, 77–80 (1999).

35. Restle, F. Theory of serial pattern learning: Structural trees. *Psychol. Rev.* **77**, 481–495 (1970).

36. Shima, K., Isoda, M., Mushiake, H. & Tanji, J. Categorization of behavioural sequences in the prefrontal cortex. *Nature* **445**, 315–318 (2006).

37. Wang, L., Uhrig, L., Jarraya, B. & Dehaene, S. Representation of numerical and sequential patterns in macaque and human brains. *Curr. Biol.* **25**, 1966–1974 (2015).

38. Ballard, I., Miller, E. M., Piantadosi, S. T., Goodman, N. D. & McClure, S. M. Beyond reward prediction errors: Human striatum updates rule values during learning. *Cereb. Cortex* **28**, 3965–3975 (2017).

39. Bhanji, J. P., Beer, J. S. & Bunge, S. A. Taking a gamble or playing by the rules: Dissociable

prefrontal systems implicated in probabilistic versus deterministic rule-based decisions. *NeuroImage* **49**, 1810–1819 (2010).

40. Naudé, J. *et al.* Nicotinic receptors in the ventral tegmental area promote uncertainty-seeking. *Nat. Neurosci.* **19**, 471–478 (2016).

41. Shanks, D. R., Wilkinson, L. & Channon, S. Relationship between priming and recognition in deterministic and probabilistic sequence learning. *J. Exp. Psychol. Learn. Mem. Cogn.* **29**, 248–261 (2003).

42. Stefaniak, N., Willems, S., Adam, S. & Meulemans, T. What is the impact of the explicit knowledge of sequence regularities on both deterministic and probabilistic serial reaction time task performance? *Mem. Cognit.* **36**, 1283–1298 (2008).

43. Téglás, E. & Bonatti, L. L. Infants anticipate probabilistic but not deterministic outcomes. *Cognition* **157**, 227–236 (2016).

44. Yeung, N. & Summerfield, C. Metacognition in human decision-making: Confidence and error monitoring. *Philos. Trans. R. Soc. B Biol. Sci.* **367**, 1310–1321 (2012).

45. Navajas, J. *et al.* The idiosyncratic nature of confidence. *Nat. Hum. Behav.* **1**, 810–818 (2017).

46. Resulaj, A., Kiani, R., Wolpert, D. M. & Shadlen, M. N. Changes of mind in decision-making. *Nature* **461**, 263–266 (2009).

47. Parpart, P., Jones, M. & Love, B. C. Heuristics as Bayesian inference under extreme priors. *Cognit. Psychol.* **102**, 127–144 (2018).

48. Song, J.-H. & Nakayama, K. Hidden cognitive states revealed in choice reaching tasks. *Trends Cogn. Sci.* **13**, 360–366 (2009).

49. Estes, W. K. Toward a statistical theory of learning. *Psychol. Rev.* **57**, 94–107 (1950).

50. Peterson, C. R. & Beach, L. R. Man as an intuitive statistician. *Psychol. Bull.* **68**, 29–46 (1967).

51. Falk, R. & Konold, C. Making sense of randomness: Implicit encoding as a basis for judgment. *Psychol. Rev.* **104**, 301 (1997).

52. Meyniel, F., Maheu, M. & Dehaene, S. Human inferences about sequences: A minimal transition probability model. *PLoS Comput. Biol.* **12**, e1005260 (2016).

53. Kounios, J. & Beeman, M. The aha! moment: The cognitive neuroscience of insight. *Curr. Dir. Psychol. Sci.* **18**, 210–216 (2009).

54. Pauli, W. M. & Jones, M. Changepoint detection versus reinforcement learning: Separable neural substrates approximate different forms of Bayesian inference. *bioRxiv* **88**, 323 (2019).

55. Wilson, R. C., Nassar, M. R. & Gold, J. I. A mixture of delta-rules approximation to bayesian inference in change-point problems. *PLoS Comput. Biol.* **9**, e1003150 (2013).

56. Huettel, S. A., Mack, P. B. & McCarthy, G. Perceiving patterns in random series: Dynamic processing of sequence in prefrontal cortex. *Nat. Neurosci.* **5**, 485–490 (2002).

57. Tervo, D. G. R. *et al.* Behavioral variability through stochastic choice and its gating by anterior cingulate cortex. *Cell* **159**, 21–32 (2014).

58. Bekinschtein, T. A. *et al.* Neural signature of the conscious processing of auditory regularities. *Proc. Natl. Acad. Sci. U. S. A.* **106**, 1672–1677 (2009).

59. Wacongne, C. *et al.* Evidence for a hierarchy of predictions and prediction errors in human cortex. *Proc. Natl. Acad. Sci.* **108**, 20754–20759 (2011).

60. Schapiro, A. C., Rogers, T. T., Cordova, N. I., Turk-Browne, N. B. & Botvinick, M. M. Neural representations of events arise from temporal community structure. *Nat. Neurosci.* **16**, 486–492 (2013).

61. Goodman, N., Tenenbaum, J., Feldman, J. & Griffiths, T. A rational analysis of rule-based concept learning. *Cogn. Sci. Multidiscip. J.* **32**, 108–154 (2008).

62. Kemp, C. & Tenenbaum, J. B. The discovery of structural form. *Proc. Natl. Acad. Sci. U. S. A.* **105**, 10687–10692 (2008).

63. Silver, D. *et al.* Mastering the game of Go with deep neural networks and tree search. *Nature* **529**, 484–489 (2016).

64. Jacobs, R. A., Jordan, M. I., Nowlan, S. J. & Hinton, G. E. Adaptive mixtures of local experts.

Neural Comput. (1991).

65. Gigerenzer, G. & Brighton, H. Homo heuristicus: Why biased minds make better inferences. *Top. Cogn. Sci.* **1**, 107–143 (2009).

66. Wolpert, D. H. The lack of a priori distinctions between learning algorithms. *Neural Comput.* (1996).

67. Radulescu, A., Niv, Y. & Ballard, I. Holistic reinforcement learning: The role of structure and attention. *Trends Cogn. Sci.* **23**, 278–292 (2019).

68. Frank, M. C. & Tenenbaum, J. B. Three ideal observer models for rule learning in simple languages. **120**, 360–371 (2011).

69. Anderson, J. R. The adaptive nature of human categorization. *Psychol. Rev.* **98**, 409–429 (1991).

70. Ashby, F. G., Alfonso-Reese, L. A., Turken, A. U. & Waldron, E. M. A neuropsychological theory of multiple systems in category learning. *Psychol. Rev.* **105**, 442–481 (1998).

71. Ashby, F. G. & Maddox, W. T. Human category learning. *Annu. Rev. Physiol.* **56**, 149–178 (2005).

72. Chomsky, N. & Halle, M. *The Sound Pattern of English*. (MIT Press, 1991).

73. Saffran, J. R. Statistical language learning. *Curr. Dir. Psychol. Sci.* **12**, 110–114 (2003).

74. Feldman, J. Minimization of Boolean complexity in human concept learning. *Nature* **407**, 630–633 (2000).

75. Gobet, F. *et al.* Chunking mechanisms in human learning. *Trends Cogn. Sci.* **5**, 236–243 (2001).

76. Planton, S. *et al.* Mental compression of binary sequences in a language of thought. *PsyArXiv* (2020) doi:10.31234/osf.io/aez4w.

77. Shannon, C. E. *A Mathematical Theory of Communication*. vol. 27 (Bell System Technical Journal, 1948).

78. Brainard, D. H. The psychophysics toolbox. *Spat. Vis.* **10**, 433–436 (1997).

79. Bayes, M. & Price, M. An Essay towards Solving a Problem in the Doctrine of Chances. By the Late Rev. Mr. Bayes, F. R. S. Communicated by Mr. Price, in a Letter to John Canton, A. M. F. R. S. *Philos. Trans. R. Soc. Lond.* **53**, 370–418 (1763).

80. Laplace, P. *Théorie Analytique des Probabilités*. (1812).

81. Gelman, A. *et al. Bayesian Data Analysis, Third Edition*. (CRC Press, 2013).

82. Tenenbaum, J. B. & Griffiths, T. L. Generalization, similarity, and Bayesian inference. *Behav. Brain Sci.* **24**, (2002).

83. Brown, P. F., Desouza, P. V., Mercer, R. L., Pietra, V. J. D. & Lai, J. C. Class-based n-gram models of natural language. *Comput. Linguist.* **18**, 467–479 (1992).

84. Jeffreys, H. *The theory of probability*. (OUP Oxford, 1998).

85. Macmillan, N. A. & Creelman, C. D. *Detection Theory: A User's Guide*. (Routledge, 2004).

86. Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D. & Iverson, G. Bayesian t-tests for accepting and rejecting the null hypothesis. *Psychon. Bull. Rev.* **16**, 225–237 (2009).

Acknowledgements

Author contributions

Additional information
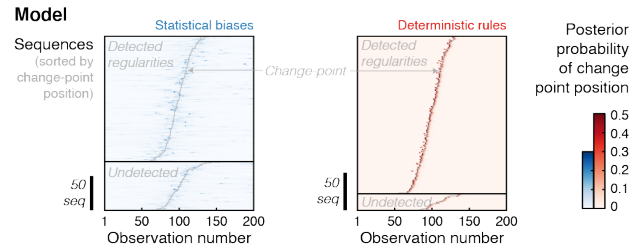
Supplementary Information is available for this paper.

Competing interests: The authors declare no competing interests.

Correspondence should be addressed to M.M.

Supplementary information

When not explicitly specify otherwise, "model" refers to the *normative two-system model*.

Supplementary figures

Supplementary Fig. 1 | Accurate inference of change-point position by the model. Posterior distribution over change-point positions marginalized over hypotheses from the model. Sequences are grouped by detection (from subjects' post-sequence reports) and sorted by change-point position (same order as *Fig. 4d,f*).
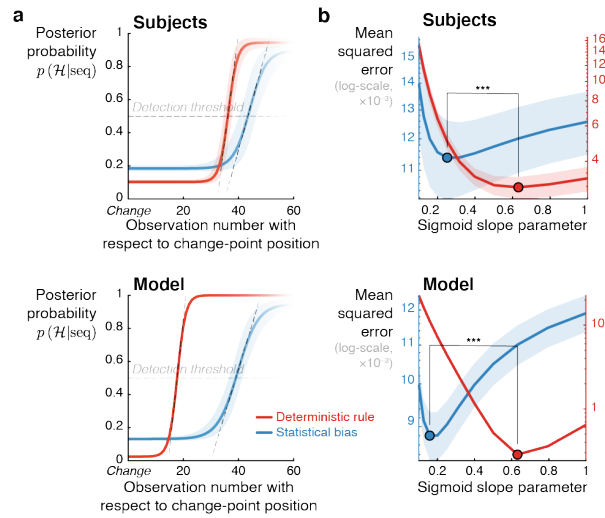
Supplementary Fig. 2 | Different detection slope for statistical biases and deterministic rules. a, *Best-fitting sigmoid functions to detection dynamics*. Sigmoid functions were fitted to single-trial dynamics of posterior probability in the non-random hypotheses. Best-fitting parameters were averaged over sequences and subjects to create these regularity-specific average sigmoid functions. **b**, *Mean-squared error over a grid of slope parameter values*. The distribution of mean squared error (MSE) is plotted as a function of a grid of values of the sigmoid slope parameter (ater minimizing the MSE over the other parameters of the sigmoid functions). In **a** and **b**, analyses were restricted to non-random sequences that were correctly identified by subjects and shaded areas correspond to the standard error of the mean computed over subjects. Stars denote significance: *** $p < 0.005$, ** $p < 0.01$, * $p < 0.05$.
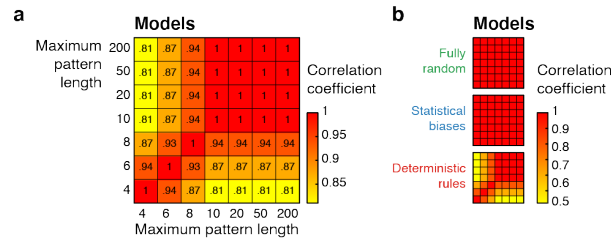
Supplementary Fig. 3 | Similarity of model inference for different maximum pattern lengths. a, *Correlation across all types of sequences*. Correlation between posterior probabilities (after conversion to cartesian coordinates to ensure an appropriate number of degrees of freedom) from different versions of the model considering different maximum lengths for pattern detection. Models that use a maximum pattern length larger than the longest patterns (i.e. 10 observations) used in the experiment make very similar inferences. **b**, *Correlation across each type of sequence*. Same as in **a** but separately for each type of sequence. The difference in inference between versions of the model using small maximum pattern lengths arise solely for sequences entailing deterministic rules, thereby suggesting that the difference originates from longest patterns remaining undetected.
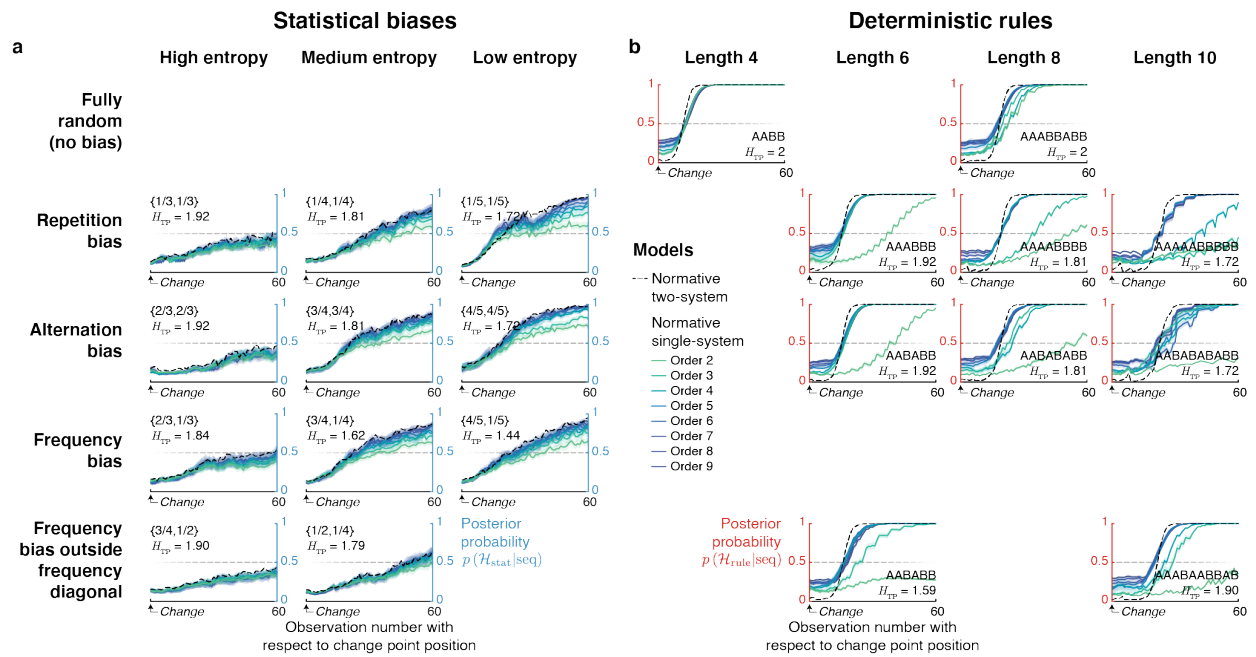
Supplementary Fig. 4 | Detection dynamics for all regularities by the *normative single-system model*. a, **b**, *In the case of statistical biases/deterministic rules*. The posterior probability of the statistical bias/deterministic rule hypothesis is displayed post change-point. Simulations are performed on sequences that were accurately identified by subjects. Error bars correspond to the standard error of the mean computed over subjects.
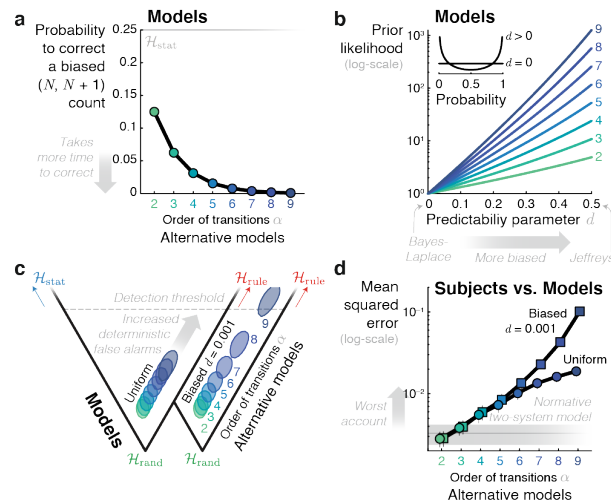
Supplementary Fig. 5 | Inflated false alarms in the *normative single-system model*. a, *Inflated false alarms follow from long-lasting counts imbalance*. In random sequences (i.e. the initial part of random-to-non-random sequences, or the entire length of the random sequences), frequency counts (i.e. number of A vs. B given one particular transition) can transiently deviate from chance level. In particular, because observations are all-or-none, frequency counts often entail a N vs. $N + 1$ imbalance. The plot shows the probability that each newly received observation would correct this light imbalance, which is nothing but the product between observing the correcting item (e.g. observing A if $N^B = N^A + 1$) and the probability of observing the particular transition for which the counts are biased. Biased counts take more time to correct (i.e. have a low correcting probability) for higher-order transitions. The posterior probability of these alternatives, with their many long-lasting imbalanced counts, thus become overall more likely. **b**, *Inflated false alarms follow from increased prior likelihood*. The posterior probability of the deterministic rule hypothesis depends upon the likelihood of the sequence under that particular hypothesis. The a priori (i.e. before receiving any observation) sequence likelihood is plotted here as a function of the strength of the prior predictability bias (which shapes the conjugate prior distribution, shown as inset) and the order of the transitions to monitor. The posterior probability of alternatives with higher-order transitions and strong predictability bias alternatives are a priori more likely. **c**, *Average position in random sequences*. Upper 10% from the triangular histograms of posterior probabilities reported in random sequences (to compare with *Fig. 6b*) are displayed for the different versions of the *normative single-system model* that monitors transition probabilities of different orders using either a uniform or biased prior. **d**, *Error of the different versions of the normative single-system model*. The mean squared error (MSE) between posterior probability in the deterministic rule hypothesis from subjects and models. A larger error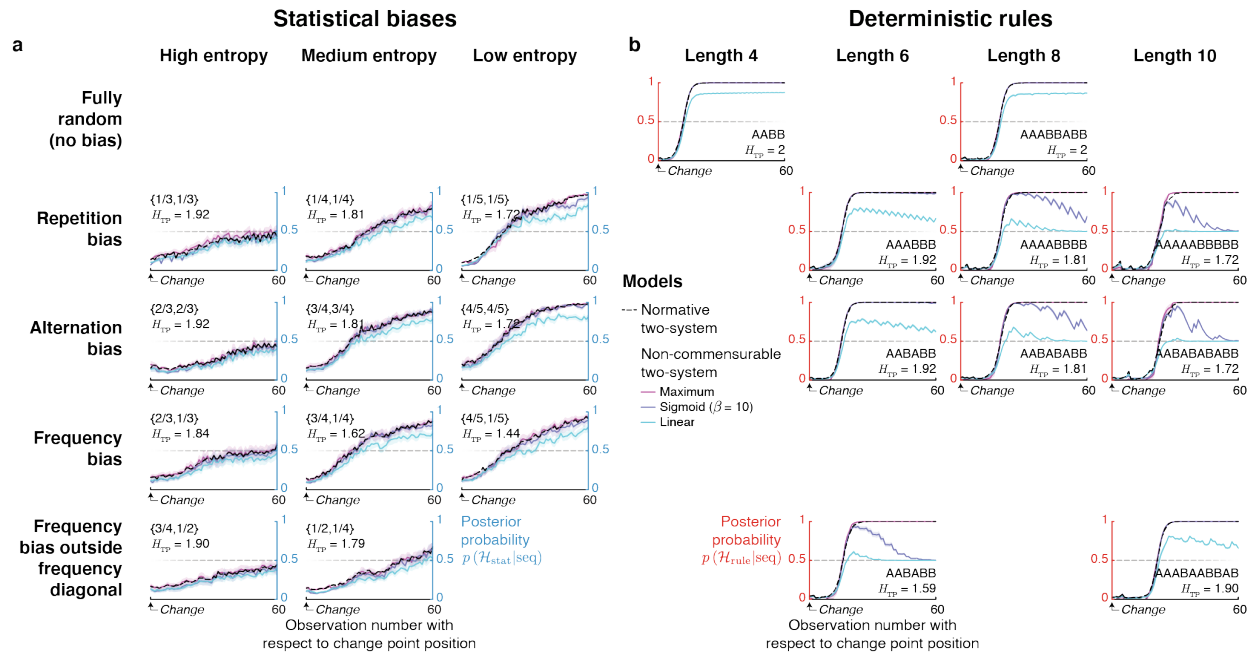 indicates a worst account of human average deterministic false alarm. Error bars correspond to the standard error of the mean computed over subjects. Stars denote significance: *** $p < 0.005$, ** $p < 0.01$, * $p < 0.05$.

**Supplementary Fig. 6 | Detection dynamics for all regularities by the *non-commensurable two-system model*. a**, **b**, *In the case of statistical biases/deterministic rules*. The pseudo posterior probability of the statistical bias/deterministic rule hypothesis is displayed post change-point. Simulations are performed on sequences that were accurately identified by subjects. Error bars correspond to the standard error of the mean computed over subjects.
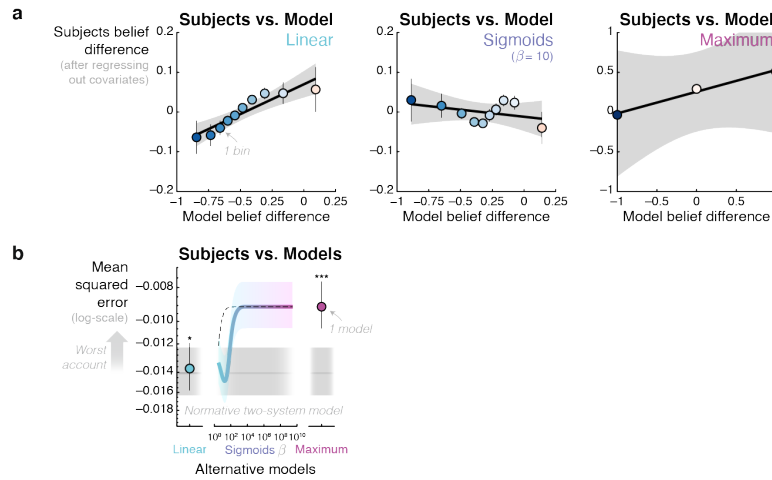
Supplementary Fig. 7 | Irrational competition between non-random hypotheses for the *non-commensurable two-system model*. a, *Belief difference from subjects as a function of the belief difference from the different versions of the non-commensurable two-system model in random sequences* (compare with *Fig. 7c*). In random sequences (i.e. the initial part of random-to-non-random sequences, or the entire length of the random sequences), the belief difference from subjects was regressed against belief difference from the different versions of the *non-commensurable two-system model*. The linear relationship between subjects and the model's belief difference is shown (in 10 bins defined using deciles) after regressing-out effect of confounding variables (i.e. the posterior probability of the random hypothesis together with its interaction with the belief difference). Error bars correspond to the standard error of the mean computed over subjects and shaded area to 95% confidence interval of the regression coefficients. **b**, *Error of the different versions of the non-commensurable two-system model*. The mean squared error (MSE) from the regression is displayed for the different versions of the *non-commensurable two-system model*. The MSE is plotted relatively to a null regression with only an offset parameter. The dashed line corresponds to versions with a sigmoid-based combination but using the log-ratio between posterior probabilities of the non-random hypotheses (instead of the difference). Error bars and shaded area correspond to the standard error of the mean computed over subjects. Stars denote significance: *** $p < 0.005$, ** $p < 0.01$, * $p < 0.05$.
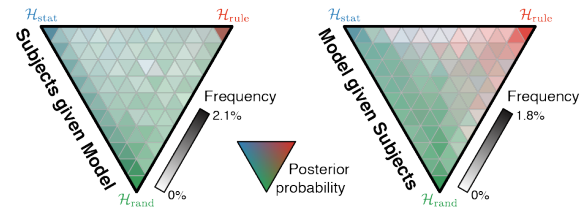
Supplementary Fig. 8 | Overall agreement between subjects and the model. The agreement map (left panel) depicts what was the hypothesis posterior probability reported by subjects (which is color-coded according to the triangular colormap) as a function of the corresponding hypothesis posterior probability from the model (averaged in 10 x 10 x 10 equally spaced bins). Transparency reflects the normalized log-number of observations in each bin. Right panel shows the same thing but for the model conditionally upon what the subjects reported. These maps are obtained by concatenating data from all subjects and all sequences together. The maps show strong agreement between subjects and the model both for extreme (corresponding colours at each summit) and intermediate (progressively changing colours in between pairs of summits) posterior probabilities.
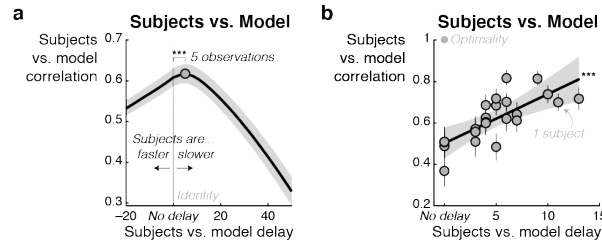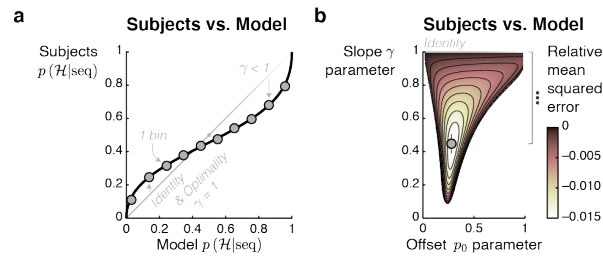
Supplementary Fig. 9 | Delay between subjects' and model's reports. **a**, *Response delay.* Coefficients from a cross-correlation between all hypothesis posterior probabilities from subjects and the model (after conversion to cartesian coordinates to ensure an appropriate number of degrees of freedom) after shifting in time one with respect to the other. The average delay corresponds to ~5 observations. Error bars correspond to the standard error of the mean computed over subjects. **b**, *Inter-subject correlation between estimated delay and correlation strength with the model.* Subjects with longer integration delay reported more accurate values. Error bars correspond to the standard error of the mean computed for each subject over sequences and the shaded area to 95% confidence interval of the regression coefficients. Stars denote significance: *** $p < 0.005$, ** $p < 0.01$, * $p < 0.05$.

1

Supplementary Fig. 10 | Non-linear mapping between subjects' and model's posterior probabilities. **a**, *Probability weighting*. All hypothesis posterior probabilities from the subjects are plotted against corresponding hypothesis posterior probabilities from the model (in 10 equally spaced bins). A nonlinear probability weighting function maps this relationship: small probabilities are overestimated by subjects and large ones are underestimated, leading to a γ slope parameter smaller than 1. Error bars correspond to the standard error of the mean computed over subjects. **b**, *Best parameter set*. Parameters of the probability weighting function (i.e. the slope γ and the offset p_0) were varied along a grid of values. Mean squared error (i.e. MSE) between transformed probabilities (with the probability weighting function) from the model and corresponding probabilities from the subjects was computed (after conversion to cartesian coordinates to ensure an appropriate number of degrees of freedom). Here, the MSE is plotted relative to the error obtained with a linear, non-weighted, mapping (i.e. $\gamma = 1$). The grey dot represents the average best parameter set (i.e. which induces the smallest error). Parameter sets for which the linear mapping provides a better account (i.e. relative MSE > 0) have been masked. Stars denote significance: *** $p < 0.005$, ** $p < 0.01$, * $p < 0.05$.

2

Supplementary note 1

Reason for the non-abrupt detection dynamics of deterministic rules in the *normative single-system model*

The *normative single-system model* predicts non-abrupt detection dynamics for deterministic rules with apparent biases in transition probabilities (*Fig. 8b*), which provided a poor account to human data (*Fig. 8c*). More specifically, detection dynamics were characterized by small steps, which occur at specific moments of pattern repetitions (*Supplementary Fig. 4*). For instance, for the AAAAABBBBB pattern, the posterior probability in the deterministic rule hypothesis (H'_{rule} monitoring higher-order transitions) suddenly increases after the A|B and B|A transitions. This is because these rare first-order transitions are surprising from the viewpoint of the statistical bias hypothesis (H_{stat} monitoring first-order transitions) whose posterior probability consequently decreases, for the benefit of the deterministic rule hypothesis. This does not happen (or only marginally) in the *normative two-system model* because, in this case, the deterministic rule hypothesis monitors the repetition of fixed patterns which permits a much more rapid increase in posterior probability, and therefore less competition with the statistical bias hypothesis.

Supplementary note 2

Additional aspects on which to reject alternative models

In the main text, the different versions of the alternative *normative single-system model* and the *non-commensurable two-system model* were rejected on the basis of peculiar detection dynamics for deterministic rules that were not observed in subjects. However, these models also fail to account for some other important aspects of human data that are observed during random sequences and which we detail below.

Normative single-system model. This alternative model detects the deterministic rules on the basis of the apparent statistical bias they induce. We have considered statistical biases corresponding to transition probabilities ranging from order 2 to 9. As the order of the transitions to monitor increased, we observed an increase in the posterior probability of the deterministic rule hypothesis during random sequences (coefficient = 0.014, CI = [0.035, 0.037], d_{Cohen} = 8.92, t_{22} = 42.8, p = 1.12 · 10^{-22}; *Supplementary Fig. 5c*). This inflated deterministic false alarm rate provided a significantly worse account of human behaviour compared to the *normative two-system model* for orders larger than 5 (difference in MSE for $p(H_{\text{rule}}|y)$ > 0.0040, d_{Cohen} > 0.60, t_{22} > 2.88, p < 0.009; *Supplementary Fig. 5d*). This is because local departures from randomness take more time to vanish (i.e. have a low correcting probability) when monitoring higher-order transitions (*Supplementary Fig. 5a*). In random sequences, the probability that each newly received observation will correct a simple $N^B = N^A + 1$ imbalance in the frequency counts is determined by:

$$p = p\left(\text{A}\right) \cdot p\left(\text{transition}\right) = \frac{1}{2} \cdot \frac{1}{2^{\alpha}}$$

We also investigated a similar version of this model but which uses a prior distribution biased for predictability and found that, even with a very weak bias (d = 0.001 instead of d = 0), the resulting inflated false alarm level became even larger compared to versions with a uniform prior

(MSE regarding $p(H_{\text{rule}}|y)$ compared to alternatives with uniform prior $> 6.59 \cdot 10^{-4}$, $d_{\text{Cohen}} > 22.8$, $t_{22} > 109.3$, $p = 1.37 \cdot 10^{-31}$; *Supplementary Fig. 5c*), and therefore also provided a significantly worse account of human behaviour compared to the *normative two-system model* for orders larger than 4 (difference in MSE for $p(H_{\text{rule}}|y) > 0.0051$, $d_{\text{Cohen}} > 0.48$, $t_{22} > 2.30$, $p < 0.031$; *Supplementary Fig. 5d*). This is mathematically caused by an increase in the prior likelihood of the sequence under the deterministic rule hypothesis, which depends upon both the order of the transitions to monitor and the strength of the prior predictability bias parameter (*Supplementary Fig. 5b*), as shown with the following mathematical rule (adapted from equation 10 in the main text):

$$p\left(y_{\text{rule}}|\mathcal{H}'_{\text{rule}}\right) = \frac{1}{2^{\alpha}} \cdot \int_{0}^{1} \prod_{i=1}^{2^{\alpha}} \text{Beta}\left(\theta_{\text{A}|T_i}|1-d, 1-d\right) \mathrm{d}\boldsymbol{\theta}$$

Non-commensurable two-system model. This alternative model considers that the posterior probability of the statistical bias and deterministic rule hypotheses cannot be expressed in the same common probabilistic currency. As a result, this model cannot rationally weigh the two non-random hypotheses against each other during random sequences. To assess how well the relative credence assigned to the non-random hypotheses by subjects related to that of the different models, we regressed subjects' difference in beliefs between non-random hypotheses against the difference in beliefs (after including confounding variables in the regression model: the posterior probability of the random hypothesis together with its interaction with the belief difference; as in the main text) from the different versions of the *non-commensurable two-system model* (*Supplementary Fig. 7a*). To compare these different versions of the model against the *normative two-system model*, we computed the mean squared error of this regression (relative to a null model with an offset predictor only). Results show that both the "linear difference" (difference in relative MSE $= 4.31 \cdot 10^{-4}$, CI $= [1.02 \cdot 10^{-4}, 7.60 \cdot 10^{-4}]$, $d_{\text{Cohen}} = 0.57$, $t_{22} = 2.71$, $p = 0.013$) and "max rule" (difference in relative MSE $= 0.0050$, CI $= [0.0029, 0.0071]$, $d_{\text{Cohen}} = 1.02$, $t_{22} = 4.87$, $p = 7.18 \cdot 10^{-}$

[5]) versions of the *non-commensurable two-system model* provided a worse account of subjects'

behaviour (*Supplementary Fig. 7b*). Versions resorting to a sigmoid function (instead of the "linear

difference" or "max rule") to map the independently-computed hypothesis posterior probabilities

(using either the difference or the log-ratio) moves the model along a continuum whose ends

correspond to the "linear difference" and "max rule" versions tested above and provided a

significantly worse account of subjects' behaviour compared to the *normative two-system model*. In

the case of a sigmoid mapping on the difference, only a small range of slope parameter (from 3.16

to 11.10) induced statistically non-significant MSE relative to the *normative two-system model* but

were associated with regression coefficients smaller than 0.

Supplementary note 3

Examples of quantitative deviations from optimality

Even though many aspects of human behaviour are well accounted for by the *normative two-system model* (see *Supplementary Fig. 8* for a summary), we are not claiming here that human behaviour is fully optimal. Instead, we observed several quantitative deviations from optimality.

Integration delay. Because of the presentation rate (one sound every 300 milliseconds), it seems a priori difficult for subjects to update immediately their reported likelihoods, thereby inducing a delay between the presentation of the observation and its effect on the reports. We estimated correlations between posterior probabilities reported by human subjects vs. the model (after conversion to cartesian coordinates to ensure an appropriate number of degrees of freedom), after shifting in time one with respect to the other, in order to estimate the most likely integration delay characterizing each subject. Integration delay was found to be significantly larger than zero (delay = 5.00 observations ≈ 1.5 second, CI = [3.49, 6.51], $d_{\text{Cohen}} = 1.43$, $t_{22} = 6.87$, $p = 6.70 \cdot 10^{-7}$; *Supplementary Fig. 9a*). Interestingly, we also found a positive correlation between the integration delay and the correlation with the model when taking into account the delay ($r_{21} = 0.74$, $p = 4.96 \cdot 10^{-5}$; *Supplementary Fig. 9b*), suggesting that more cautious subjects reported more accurate values.

Probability weighting. The literature on probabilistic inference frequently report that humans overestimate small probabilities and underestimate large probabilities (*Supplementary Fig. 10a*), a phenomenon termed probability weighting and which is commonly described by the following function which maps objective (p) to subjective (p^*) probabilities:

$$p^* = \frac{1}{1 + \exp\left(-\gamma \cdot \log \frac{p}{1-p} - (1 - \gamma) \cdot \log \frac{p_0}{1-p_0}\right)}$$

To assess and quantify the distortion (relative to optimality) of subjects' reported posterior probability in the different hypotheses, we estimated the mean squared error (MSE, relative to a linear mapping) between posterior probabilities reported by human subjects vs. the model (using cartesian coordinates and not barycentric ones to use an appropriate number of degrees of freedom). More precisely, we used a grid from 0.01 and 1.99 for the slope parameter (γ) and from 0.01 to 0.99 for the offset parameter (p_0), with 0.01 steps in both cases, and found the set of parameters inducing the smallest MSE. We found the classic under-/overestimation effect in our data: the slope parameter was found to be significantly smaller than 1 (difference of γ slope parameter to 1 = –0.55, CI = [–0.67, –0.43], d_{Cohen} = –2.01, t_{22} = –9.65, $p = 2.29 \cdot 10^{-9}$; *Supplementary Fig. 10b*) and induced a MSE significantly smaller that an unbiased linear mapping (difference in MSE = –0.078, CI = [–0.031, –0.015], d_{Cohen} = –1.21, t_{22} = –5.82, $p = 7.38 \cdot 10^{-6}$).